



HAL
open science

Image processing algorithms for the visualization of interventional devices in X-ray fluoroscopy

Vincent Bismuth

► **To cite this version:**

Vincent Bismuth. Image processing algorithms for the visualization of interventional devices in X-ray fluoroscopy. Other [cs.OH]. Université Paris-Est, 2012. English. NNT : 2012PEST1062 . tel-00747682v2

HAL Id: tel-00747682

<https://theses.hal.science/tel-00747682v2>

Submitted on 8 Feb 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITE PARIS EST



Ecole doctorale MSTIC

Image processing algorithms for the visualization of interventional devices in X-ray fluoroscopy

Algorithmes de traitement d'images pour la visualisation
d'outils interventionnels dans des séquences de
fluoroscopie par rayons X.

par

Vincent Bismuth

Thèse soumise pour l'obtention du grade de docteur en informatique
de l'université Paris Est

au

laboratoire d'informatique Gaspard-Monge
ESIEE

Soutenue le 9 janvier 2012

Jury

Isabelle Bloch - Présidente
Grégoire Malandain - Rapporteur
Nassir Navab - Rapporteur
François Funck - Examineur
Hans Reiber - Examineur
Régis Vaillant - Examineur
Laurent Najman - Directeur de thèse

*“Caminante, son tus huellas
el camino y nada más;
Caminante, no hay camino,
se hace camino al andar.
Al andar se hace el camino,
y al volver la vista atrás
se ve la senda que nunca
se ha de volver a pisar.
Caminante no hay camino
sino estelas en la mar.”*

Antonio Machado - Extracto de Proverbios y cantares (XXIX)

L'université n'entend donner aucune approbation ni improbation aux opinions émises dans les thèses : ces opinions doivent être considérées comme propres à leurs auteurs.

UNIVERSITE PARIS EST

Abstract

laboratoire d'informatique Gaspard-Monge
ESIEE

Doctor of Philosophy

**Image processing algorithms for the visualization of interventional devices
in X-ray fluoroscopy.**

by Vincent Bismuth

Stent implantation is the most common treatment of coronary heart disease, one of the major causes of death worldwide. During a stenting procedure, the clinician inserts interventional devices inside the patient's vasculature. The navigation of the devices inside the patient's anatomy is monitored in real-time, under X-ray fluoroscopy. Three specific interventional devices play a key role in this procedure: the guide-wire, the angioplasty balloon and the stent. The guide-wire appears in the images as a thin curvilinear structure. The angioplasty balloon, that has two characteristic marker-balls at its extremities, is mounted on the guide-wire. The stent is a 3D metallic mesh, whose appearance is complex in the fluoroscopic images. Stents are barely visible, but the proper assessment of their deployment is key to the procedure.

The objective of the work presented in this thesis is twofold. On the first hand, we aim at designing, studying and validating image processing techniques that improve the visualization of stents. On the second hand, we study the processing of curvilinear structures (like guide-wires) for which we propose a new image processing technique. We present algorithms dedicated to the 2D and 3D visualization of stents. Since the stent is hardly visible, we do not intend to directly locate it by image processing means in the images. The position and motion of the stent are inferred from the location of two landmarks: the angioplasty balloon and the guide-wire, which have characteristic shapes. To this aim, we perform automated detection, tracking and registration of these landmarks.

The cornerstone of our 2D stent visualization enhancement technique is the use of the landmarks to perform motion compensated noise reduction. We evaluated the performance of this technique for 2D stent visualization over a large database of clinical data (nearly 200 cases). The results demonstrate that our method outperforms previous state of the art techniques in terms of image quality. A comprehensive validation confirmed that we reached the level of performance required for the commercial introduction of

our algorithm. It is currently deployed in a large number of clinical sites worldwide. The 3D stent visualization that we propose, uses the landmarks to achieve motion compensated tomographic reconstruction. We show preliminary results over 22 clinical cases. Our method seems to outperform previous state of the art techniques both in terms of automation and image quality.

The previous stent visualization methods involve the segmentation of the part of the guide-wire extending through the stent. We propose a generic tool, called the Polygonal Path Image (PPI), to process such curvilinear structures. The PPI relies on the concept of locally optimal paths. One of its main advantages is that it unifies the concepts of several previous state of the art techniques in a single formalism. Moreover the PPI enables to control the smoothness and the length of the structures to segment. Its parameterization is simple and intuitive. In order to fully benefit from the PPI, we propose an efficient scheme to compute it. We demonstrate its applicability for the task of automated guide-wire segmentation, for which it outperforms previous state of the art techniques.

Keywords: stent, guide-wire, cardiology, image processing, Xray.

UNIVERSITE PARIS EST

Résumé

laboratoire d'informatique Gaspard-Monge
ESIEE

Doctorat

Algorithmes de traitement d'images pour la visualisation d'outils interventionnels dans des séquences de fluoroscopie par rayons X.

par [Vincent Bismuth](#)

La pose de stent est l'option de traitement la plus courante de la maladie coronarienne, l'une des principales causes de mortalité dans le monde. Lors d'une procédure de pose de stent, le médecin insère des outils chirurgicaux dans le réseau vasculaire du patient. La progression de ces outils à l'intérieur du corps est suivie en temps réel sous uroscopie par rayons X. Trois outils, en particulier, jouent un rôle crucial dans la procédure : le guide, le ballon d'angioplastie et le stent. Le guide apparaît dans les images sous la forme d'une structure curviligne. Le ballon, monté sur le guide, est équipé de deux marqueurs radio-opaques à ses extrémités. Le stent est un maillage métallique qui se projette en une forme complexe dans les images uroscopique. Le stent, dont le bon déploiement est essentiel au succès du geste médical, est souvent très difficilement visible dans les images.

Les travaux présentés dans cette thèse poursuivent un double objectif. Il s'agit d'une part, de concevoir, d'étudier et de valider des techniques de traitement d'image visant à améliorer la visualisation des stents. D'autre part, nous étudions le traitement des structures curvilignes (comme les guides) pour lesquelles nous proposons un nouvel outil. Nous présentons des algorithmes de traitement d'image dédiés à la visualisation 2D et 3D des stents. Nous sommes amenés, dans ce but, à détecter, suivre et recalculer, de manière complètement automatique, les outils nécessaires à la pose de stent que sont le guide et le ballon.

Le stent étant à peine visible dans les images, nous ne cherchons pas à le localiser directement à l'aide de techniques de traitement d'images. La position et le mouvement du stent sont déterminées par nos algorithmes à partir de celles de deux amers: le guide et le ballon qui ont des formes caractéristiques. Nous effectuons donc, dans ce but, la détection, le suivi et le recalage de ces amers, et ce de manière complément automatique. Le cœur de notre méthode de visualisation des stents en 2D réside dans l'utilisation des

amers pour effectuer un débruitage compensé en mouvement. Nous avons évalué la performance de ces outils pour la visualisation des stents en 2D, sur une large base de près de 200 cas cliniques. Il en ressort que notre méthode surpasse les méthodes utilisées jusqu'ici sur le plan de la qualité image. La validation exhaustive que nous avons menée, a confirmé que nous avons atteint un niveau compatible avec son introduction commerciale.

Le logiciel qui en résulte est désormais installé sur un grand nombre de sites cliniques, où il est régulièrement utilisé. La méthode de visualisation 3D des stents que nous proposons utilise les amers pour effectuer une reconstruction tomographique compensée en mouvement. Nous exposons des résultats préliminaires sur une base de 22 cas cliniques. Il semble que notre méthode surpasse les méthodes précédemment employées aussi bien du point de vue de la qualité image que de l'automatisation.

Les méthodes de visualisation des stents que nous proposons s'appuient sur la segmentation de la portion du guide qui traverse le stent. Nous proposons un nouvel outil pour le traitement de telles structures curvilignes que nous appelons : l'Image de Chemins Polygonaux (acronyme PPI en anglais). Cet outil repose sur la notion de chemin localement optimal. L'un des principaux avantages du PPI est d'unifier dans un même cadre différents concepts pré-existants. De plus, il permet de contrôler la régularité et la longueur des structures à traiter avec une paramétrisation simple et intuitive. Afin de tirer pleinement parti des performances du PPI nous proposons un schéma algorithmique efficace pour le calculer. Nous illustrons ses utilisations pour la segmentation automatique de guide ou il surpasse les techniques existantes.

Mots-clefs: stent, guide, cardiologie, traitement d'image, rayons X.

Acknowledgements

I want to thank:

Isabelle Bloch who has been president of the jury which was both an honor and a motivation for me.

The “rapporteurs” of this thesis: Nassir Navab and Grégoire Malandain for kindly accepting this demanding role. Living up to your expectations was a strong motivation to keep improving my work, especially during the last and most tiring phases of writing the manuscript.

Régis Vaillant and Laurent Najman: The list of the reasons why I would like to thank both of you is too long to be written here, and can hardly be expressed with words. I am just happy to have the opportunity to write these words for you here: Thank you.

Dr Funck for his support, explanations, deep implication in the StentViz study and for being part of the jury. I take advantage of this space to thank all the team of the interventional cardiology department of “centre hospitalier René Dubos” for their commitment to the StentViz study and for their kindness.

Hans Reiber, for contributing to the jury and bringing both an academic and industrial perspective to my work.

I would like to dedicate a special note here to the memory of Pr. Til Aach who had accepted to be part of the jury but whose presence was made impossible. We deplore his sudden passing away.

Sylvie Cach for her support and recommendations throughout these three years.
Francois Kotian for understanding, supporting and making my PhD project possible inside GE.

Guillaume Buc for continuing the project.

My colleagues: The members of my team for encouraging me on this path and for our daily scientific and non scientific exchanges. All the colleagues who are also my friends and that I enjoy meeting everyday. In the context of this work, I am particularly grateful to Cyril Riddell and Dave Langan for helping me improve the redaction of my publication in MedIA. I want to express a special “Thanks!” to Sebastien Gorges for his wise advices during our weekly, unofficial, PhD meetings.

The interns I coached who contributed to this work: Working with you all was a mutual enrichment. Among them, a special thanks to Xi Yue, Paul Delafon and Angela Chieh whose contributions are more directly linked to some of the work reported here.

The ESIEE team for hosting me during these three years. I am especially grateful to

Hugues Talbot for his insights on curvilinear structure processing and his constant involvement in this part of my work.

My parents, brother and sister, family and friends for their love and support. Special thanks to my close family for what they endured during the last stages of the redaction of this manuscript (including proof-reading some pieces of it). Among my friends, a special thought for John who, beyond his unfailing support, has been an inspiring model of determination.

All the people I met, who gave me the desire to get interest in maths, program computers and appreciate the beauty of these arts - including some of my uncles: Philippe, Laurent and Pierre.

All the attendees of my PhD defense, who supported me for this final effort by being there that day, knowing that some of you came a long way, and others were not really delighted with a two hour talk in English. . .

The martial art group I belong to. A special thanks to my teachers, in chronological order: Michel Casteran, Christian Derval and Patrick Bittan. It may sound odd here for those who are not into this path, but they are also a strong part of this accomplishment. The great composers and interprets of MPB. I have been listening to this music so much while working on my PhD. . . It looks like it is the theme song of so many great moments and meetings in my life.

Abbreviations

AMI	A cute M ycocardial I nfarction
AUR	A rea U nder C urve
BMS	B are M etal S tent
CTO	C hronic T otal O clusion
CNR	C ontrast to N oise R atio
DES	D rug E luting S tent
DSA	D igital S ubtracted A ngiography
DSE	D igital S tent E nhancement
FFR	F ractional F low R eserve
fps	f rames p er s econd
IVUS	I ntra V ascular U ltra S ound
MI	M ycocardial I nfarction
MTF	M ean T ransfer F unction
OCT	O ptical C oherence T omography
PCI	P ercutaneous C oronary I ntervention
QCA	Q uantitative C oronary A nalysis
RLT	R otated L ine T emplate
ROI	R egion O f I nterest
SNR	S ignal to N oise R atio
SOD	S econd O rders D erivative
STEMI	S T E levation M ycocardial I nfarction

Contents

Abstract	iii
Résumé	v
Acknowledgements	vii
Abbreviations	ix
List of Figures	xv
List of Tables	xxviii
1 Introduction	1
1.1 Pathology and treatment	1
1.1.1 The cardiovascular system	2
1.1.2 The atheromatous stenosis and the myocardial infarction	3
1.1.3 Treatment options	4
1.1.4 Drug administration	4
1.1.5 Angioplasty	6
1.1.6 Coronary artery bypass graft (CABG)	6
1.1.7 Risk factor Management	6
1.2 A more detailed view on stenting	7
1.2.1 A short presentation of interventional cardiology	8
1.2.2 The cathlab	9
1.2.3 Breakdown of a stenting procedure	11
1.2.4 On the importance of assessing stent deployment and apposition	12
1.2.5 A word about calcifications	13
1.2.6 Synthesis of the needs of the clinician in PCI	14
1.3 Stent visualization in the cathlab	14
1.4 Technical context	16
1.4.1 Interventional device detection and tracking	16
1.4.2 Digital stent enhancement	16
1.4.3 Context of the thesis	17
1.5 Structure of this thesis	18
1.6 Main technical contributions	20
1.6.1 Stent visualization	20

1.6.2	Interventional tool detection and tracking	21
1.6.3	Curvilinear structure enhancement	24
1.6.4	Methodological contributions	25
1.6.5	Material published for the first time	25
1.7	The future of the cathlab in interventional cardiology	26
1.8	Conclusion	28
2	A stent visualization enhancement algorithm	29
2.1	Abstract	29
2.2	Introduction	30
2.3	Method	33
2.3.1	Clinicians' needs	33
2.3.2	From problem statement to algorithm design	34
2.3.2.1	Contrast	34
2.3.2.2	Zoom	35
2.3.2.3	Sharpness	35
2.3.2.4	Noise	36
2.3.2.5	DSE algorithm design take away	44
2.3.3	DSE algorithm	48
2.3.3.1	Algorithm overview	48
2.3.3.2	Guide-wire segmentation	48
2.3.3.3	Constrained registration derived from a stent deformation model	51
2.3.3.4	Non-linear Registration	53
2.3.4	Settings and algorithm validation	55
2.4	Results	55
2.4.1	Landmark detection and localization performance	56
2.4.1.1	Accuracy and robustness evaluation strategy	56
2.4.1.2	Marker ball detection robustness	56
2.4.1.3	Marker ball segmentation accuracy	57
2.4.1.4	Guide-wire segmentation accuracy and robustness	57
2.4.2	Image quality improvement	59
2.4.2.1	Data acquisition	60
2.4.2.2	Image quality and clinical benefit results	62
2.4.3	Non linear versus linear registration	64
2.4.3.1	Quantitative comparison of the transforms	64
2.4.3.2	Impact of the transform on image quality	66
2.4.4	Execution time	69
2.4.5	Comparison with existing software	69
2.4.5.1	Data collected and stent information	69
2.4.5.2	Comparison method	70
2.4.5.3	Results of the comparison	70
2.5	Discussion	71
2.5.1	Non linear versus linear registration	71
2.5.2	Comparison with related work	73
2.5.2.1	Comparison of the technical aspects	73
2.5.2.2	Comparison of the improvement in image quality	74

2.5.2.3	On the clinical usefulness of DSE	75
2.5.3	The role of DSE in clinical practice	75
2.5.4	Selected cases	76
2.5.5	Hardware vs software DSE	77
2.5.6	Observed limitations of the technique	77
2.6	Conclusion	80
3	Subtraction of the guide-wire in enhanced stent images	81
3.1	Abstract	81
3.2	Introduction	82
3.3	Background	82
3.3.1	The transparent layer model for Xray images	83
3.3.2	Layer separation in transparent images	85
3.4	Method	87
3.4.1	Layer separation in DSE images	88
3.4.1.1	Background layer separation	88
3.4.1.2	Guide-wire layer separation	88
3.4.2	Simultaneous display of two DSE images	90
3.4.2.1	General technique	90
3.4.2.2	Implementation details	92
3.5	Results	93
3.6	Discussion	94
3.6.1	The benefits of guide-wire subtraction	94
3.6.2	Position of the guide-wire subtraction in our processing chain	95
3.6.3	Limitations	96
3.7	Conclusion	96
4	3D stent reconstruction	101
4.1	Abstract	101
4.2	Background	101
4.2.1	Clinical Background	101
4.2.2	State of the art	102
4.2.2.1	3D stent reconstruction	102
4.2.2.2	Marker detection and tracking	104
4.2.2.3	Overview of the chapter	104
4.2.3	The necessity of a new algorithm	105
4.3	Method	106
4.3.1	Clinical sequence database	106
4.3.2	Marker detection and tracking in rotational acquisitions	106
4.3.2.1	High-level architecture of the algorithm	106
4.3.2.2	Pre-processing	107
4.3.2.3	Pair detection	107
4.3.2.4	Short track construction	108
4.3.2.5	Track merging	108
	Presentation of the algorithm	108
	Algorithm	109
	Track continuity score	109

4.3.2.6	Determination of the marker track and refinement of the track	111
	False positive reduction	111
	False negative reduction	113
4.3.2.7	Track refinement	115
4.3.2.8	Labeling of the markers	115
4.3.2.9	3D reconstruction of the marker line segment	115
4.3.2.10	Application to labeling and track refinement	117
4.3.3	Guide-wire subtraction	117
4.3.4	On the choice of the registration for 3D stent reconstruction	118
4.3.4.1	Linear constant radius registration	120
4.3.4.2	Guide-wire based registration	122
4.4	Results	123
4.4.1	Marker detection and tracking	123
4.4.2	Detailed results of the different algorithms	124
4.5	3D stent Reconstructions	125
4.5.1	Guide-wire subtraction and registration	128
4.6	Discussion	130
4.6.1	Comparison with state of the art	130
4.6.2	On the acquisition	131
4.6.3	Future work	133
4.7	Conclusion	133
5	Curvilinear structure enhancement	134
5.1	Abstract	134
5.2	Introduction	134
5.3	Background	135
5.3.1	State of the art	135
5.3.2	Guide-wire detection in X-ray fluoroscopy	137
5.3.2.1	Motivations	137
5.3.2.2	The technical challenge	138
5.3.3	Evaluating and comparing line enhancement methods	138
5.3.3.1	About ROC curves	139
5.3.3.2	Application to guide-wire detection	142
5.3.3.3	Output of previous line enhancement technique comparisons	145
5.3.3.4	Selected line enhancement methods for comparison	147
5.4	The polygonal path image	147
5.4.1	Motivations	147
5.4.2	The polygonal path image: \mathfrak{P}	149
5.4.2.1	Locally optimal paths	149
5.4.3	Structure of the path image	152
5.4.4	Efficient algorithm for the computation of \mathfrak{P}	153
5.4.5	Curvilinear structure enhancement techniques derived from the polygonal path image	157
5.4.5.1	The cost image	157
5.4.5.2	Path voting	158

5.4.5.3	Pruning	162
5.4.6	Selecting path according to smoothness	163
5.4.7	Estimating local orientation	164
5.4.8	Links with other existing techniques	166
5.4.8.1	Geodesic shortest path	166
5.4.8.2	Geodesic voting	169
5.4.8.3	Hough transform	172
5.5	Results	172
5.5.1	Database	172
5.5.2	Setting of the techniques	172
5.5.3	Potential cost image	173
5.5.4	Qualitative results	173
5.5.5	Quantitative results	174
5.6	Conclusion and further work	179
6	Conclusion and perspectives	185
6.1	Main results	186
6.1.1	2D and 3D stent visualization enhancement	186
6.1.2	Image processing tools for interventional imaging	186
6.1.3	Curvilinear structure processing	187
6.1.4	Methodology	187
6.2	Limitations and technical perspectives	188
6.2.1	Limitations	188
6.2.1.1	Image quality of enhanced stent images	188
6.2.1.2	Applicability of stent visualization techniques	188
6.2.1.3	Co-visualization of stents and vessels	189
6.2.2	Perspectives	189
6.2.2.1	Layer separation	189
6.2.2.2	Curvilinear structure processing	189
6.2.2.3	Interventional tool detection and tracking	189
6.2.2.4	Clinical impact and applications	190
A	Publication list	191
A.1	Journal	191
A.2	Conferences	191
A.3	Patents	192
A.4	Supervised master's thesis	192
B	Resume	193
	Bibliography	195

List of Figures

1.1	Simplified view of the cardiovascular system (Copyright © Addison Wesley Longman, Inc publishing as Benjamin Cummings).	2
1.2	Overview of the coronary arteries [LtOH10].	3
1.3	Progression of the atheromatous stenosis over decades, potentially rupturing putting the patient at risk of major adverse events (image from [Chi06]).	5
1.4	Visualization of a diseased artery under Xray fluoroscopy (angiography) with an Innova 2100 Xray system (GE)	5
1.5	Quantitative coronary analysis allows the clinician to estimate the percentage of lumen reduction induced by the presence of the stenosis (Quantitative Analysis Software developed by Pie Medical Imaging BV ©; pictures courtesy by Pie Medical Imaging BV. Reproduced from [esa11]).	6
1.6	Balloon angioplasty [NIH10b] and stenting [NIH10c] procedures.	7
1.7	Illustration of a bypass [NIH10a].	7
1.8	Angiograms of Dotter’s 1st catheter patient, Laura Shaw: A) before transluminal dilation of the left superficial femoral artery, B) immediately after dilation, and C) 3 weeks after the procedure (image from [Pay01]).	9
1.9	Visualization of a right coronary artery under before and during contrast media injection. We note that in the absence of contrast the medical tools inside the artery are visible but not the artery itself. When contrast media is injected, the artery is visible but the medical tools become hardly visible (images obtained on an Innova 2100 system of GE healthcare).	9
1.10	Various popular coronary stents. From left to right: Cypher [Hos07], Endeavor [ptc09], Taxus Liberte [DC09], Xience V [DC10b]. Note that the metal meshes are quite different from one stent to another and can be rather complex.	10
1.11	In stent restenosis (image from [NIH10d]). This complication was common with bare metal stents. It decreased very significantly with the introduction of drug eluting stents.	10
1.12	Innova system - GE Healthcare	11
1.13	Calcification are very well visualized with Computed Tomography [DC10a]. They appear in white, around the lumen of the artery.	14
1.14	Stent cross section in OCT versus conventional angiography [Lod10].	15
1.15	Coronary stent image with conventional Xray (image zoomed on the stent). Note that the stent is almost invisible, only the guide-wire and the two are clearly visible (image obtained on an Innova 2100 system from GE).	16

1.16	Clinicians performing an angioplasty. Note the use of the stent visualization enhancement software on the top left screen that is the output of the work presented in Chapter 2. This image depicts a stent that has just been deployed with an image quality far superior to conventional Xray [FGBV09]. Image courtesy of Dr. T. Lefèvre, Dr. M-C. Morice, Dr. T. Hovasse, Dr. B. Chevalier, Dr. Y. Louvard - Institut Cardiovasculaire Paris Sud, Massy, France	19
1.17	Stent visualization enhancement results obtained with: (a) state of the art linear registration, (b) our non-linear one. Observe that the stent is sharper in (b).	20
1.18	Planar layer modeling of a DSE image: 3 separate layers, namely the anatomical background, the guide-wire and the stent, are combined in the image formation mechanism.	21
1.19	Left: input image. Center: DSE image (note that the background layer has been removed). Right: guide-wire subtracted image.	21
1.20	3D coronary stent reconstruction at a bifurcation.	22
1.21	A 3D stent slice: on the left classical processing, on the right the guide-wire has been subtracted between the markers. The stent border that is hidden by the guide-wire on the left is visible on the right, at the location indicated by the arrows.	22
1.22	Volume rendering of a reconstructed stent. Left: state of the art reconstruction technique. Right: The reconstruction benefited from the extension of our stent techniques to the 3D case. The motion compensation is performed with non-linear registration and the guide-wire has been removed from the image. Notice that the overall stent shape is easier to delineate.	22
1.23	Left: a frame of a short static acquisition. Center: automatically segmented markers and guide-wire. Right: Close up on the segmented interventional tools.	23
1.24	Markers and guide-wire automatically segmented in three frames of a rotational acquisition (top row: original images - bottom row: segmented interventional tools). Observe that the markers exchange their relative positions as a result of the gantry rotation.	23
1.25	Top left: an extract from a cine image depicting a set of smooth paths between two endpoints. Top center: A clinical image corrupted with a high level of noise. Top right: guide-wire segmented with our method between the two markerballs. Bottom left: A clinical image depicting a guide-wire. Bottom center: 3 candidate curves (note that the blue one is well aligned with the guide-wire centerline whereas the red and green ones are not). Bottom right: The three averaged profiles (same color code as bottom center). Abscissa is the signed distance to the curve and ordinate is the pixel intensity. Only the profile of the blue curve that is aligned with the guide-wire exhibits a marked minimum at its center.	24
1.26	Illustration of two applications of \mathfrak{P} . Top left: a clinical image illustrating a guide-wire. Top right: line enhancement by smooth local path voting. Bottom left: aggressive contrast setting on the image displayed in top right reveals its structure: A path attracted by linear structures originates from each pixel of the image. Bottom right: Direction field estimated with \mathfrak{P}	26

1.27	Top [Med11] principle of the FFR: the device allows to compute the ratio of pressures distally and proximally to the lesion in the hyperemic state. Bottom [CJHJ05] effect of thrombus aspiration and balloon dilatation in an occluded artery.	27
2.1	Left: a coronary stent imaged with Xray. Right: Result after DSE.	29
2.2	(a) An image from a clinical image sequence with the clinician's tools positioned in the patient's vasculature; (b) the region in (a) delimited by the gray square with annotation of tools/devices; (c) the region in (b) inside the gray square with a deployed stent; (d) demonstration of the improved stent visualization by the application of DSE to (a). This last image has been produced with StentViz, the DSE feature of GE, that is partly described in this chapter.	31
2.3	General architecture of a DSE algorithm.	32
2.4	Left: Illustration of a deployed stent in an x-ray image processed with standard contrast enhancement techniques, which may be considered a form of DSE. The resulting stent visualization is substantially inferior to that produced with a dedicated technique, as shown on the right.	35
2.5	Study of the impact of the misregistration error. Image after the application of DSE with various misregistration magnitudes. The images were produced by registering thirty frames with a registration error following a uniform distribution and averaging the frames. The misregistration magnitude is indicated in pixel units below each image.	37
2.6	Schematic illustration of the limitations of the linear, marker-ball based registration realized with an image processing software for explanatory purpose (not necessarily geometrically exact). The two top images are two images of a same stent. Bottom is the marker-ball based registration. Although the two marker-balls are accurately registered, the stent, that underwent a non-linear deformation is not accurately registered. The guide-wire that supports the marker-balls is also undergoing a non-linear deformation. We use it as additional landmark in our <i>nonlinear</i> registration. The stent image is derived from a photo of a Cypher coronary stent [tim05]	37
2.7	Standard deviation of the estimators versus the number of images in the input sequence. Results obtained with one billion simulations. Observe the general decreased in $n^{-1/2}$ predicted by the theory for the empirical mean, the trimmed means and the median.	40
2.8	Application of the background subtraction to a clinical image. Left: Original. Right: Background subtracted. Observe that most of the structures have disappeared.	41
2.9	Simulated vs estimated noise on two clinical cases for the 9 estimators: red and purple are estimated, blue and green are simulated. The overall precision of our model is typically 10%.	41
2.10	Result of the different estimators on a clinical case. Left column: result obtained with \bar{X}_n (reference). Center: Result obtained with the considered estimator. Right: difference image. From top to bottom: $X_{(0)}$, $X_{([n/4])}$, $X_{([n/2])}$, $X_{([3n/4])}$, $X_{(n)}$, $\bar{T}_{n, [\frac{p}{2}n], [\frac{p}{2}n]}$, $\bar{T}_{n, [pn], 0}$, $\bar{T}_{n, 0, [pn]}$	45

2.11	Combined image. From left to right: \bar{X}_n , $\bar{T}_{n, [\frac{p}{2}n], [\frac{p}{2}n]}$, $\bar{T}_{n, [pn], 0}$ and $\bar{T}_{n, 0, [pn]}$. Observe that $\bar{T}_{n, [\frac{p}{2}n], [\frac{p}{2}n]}$ and $\bar{T}_{n, [pn], 0}$ allow to attenuate the undesirable guide-wire multiplication effect far from the landmarks.	46
2.12	Noise reduction as a function of the number of combined frames (1, 10, 20, 30, 40, 50 and 60 when available) with the empirical mean estimator for four different clinical sequences (one color for each: red, blue, purple and green). Solid line, manual estimation, dashed line theoretical noise with the $\sigma n^{-1/2}$ model. Ordinate: noise standard deviation; Abscissa: number of averaged frames.	46
2.13	DSE images for different numbers of averaged frames: from top to bottom 1, 10, 20, 30, 40, 50 and (when available) 60. Image have undergone background subtraction as recommended in Section 2.3.2.1. Observe on the very low stent visibility in the original data on top row. Visibility increases rapidly in the first rows, then stagnates.	47
2.14	Block diagram of our DSE technique.	49
2.15	(a) Original image, (b) top hat transform applied to (a).	49
2.16	Illustration of the “detection and tracking” brick of our DSE. Top row, from left to right: original image, candidate points, candidate pairs formed out of the points. Bottom row, from left to right: three successive frames and the detected candidate pairs. Among them the tracking was able to build five relevant tracks.	49
2.17	(a) A set of tested parametric curves to segment a guide-wire between two markers. (b) In this clinical image, the dark line along the guide-wire depicts a parametric curve to be tested. The set of dark short segments represent the perpendicular profiles that are computed for this curve. (c) Illustration of the values along one of the profiles of image (b). We see that the noise is high and there is no particular pattern in this curve. (d) Averaged profile. We see that the curve follows the expected pattern: a smooth curve with a significant minimum at its center.	51
2.18	Top left: an extract from a cine image depicting the guide-wire, the stent and the markerballs. Top right: 3 candidate curves (note that the blue one is well aligned with the guide-wire centerline whereas the red and green ones are not). Center: All the profiles extracted for each curve (same color code as top right). On the abscissa is the signed distance to the curve and on the ordinate the pixel intensity. Bottom: The three average profiles. Only the one of the blue curve that is aligned with the guide-wire exhibits a marked minimum at its center.	52
2.19	(a) Notations for the curve based registration. (b) Schematic illustration of the motion field on the landmarks	54
2.20	Coronary stents are very rigid on their radial axis to handle the pressure of the artery wall, but very flexible in the perpendicular direction to conform to the curvature of the artery.	55
2.21	In dark (top), histogram of the percentage of the true positive per image sequence. In gray, histogram of false positive per image sequence (center). In white (bottom), histogram of false negative per image sequence.	58
2.22	An image of a synthetic marker-ball and guide-wire used to quantify the precision of the marker-ball segmentation.	59

2.23	Figures illustrating the characterization of the guide-wire segmentation. (a) Input clinical image, (b) Guide wire segmented on image (a), (c) Image (a) degraded up to three times the original noise, (d) Guide wire segmented on image (c), (e) Another input clinical image, (f) Ground truth (in black) and actual segmentation (in white) of the guide-wire in image (e) degraded. The Hausdorff distance, of 12 pixels in this case, is illustrated by the dashed white line.	60
2.24	Histogram of the Hausdorff distances in pixels between the ground truth guide wires and the ones segmented on degraded images.	61
2.25	Image quality scores histograms. In gray the values before DSE and in black after.	64
2.26	Image quality score histograms in fluoro and cine separately before and after DSE (StentViz).	64
2.27	Histogram of the DSE score in cine minus the score in fluoro averaged over the two raters.	65
2.28	(a) Image obtained with StentViz; (b) I_{Δ} ; (c) ROI to compute Δ . Observe that on the markers the two motion fields are identical (black spots in (b)) but differ on the rest of the image. (d) All the segmented guide wires registered by the marker ball based techniques. Note that they are not accurately registered outside of the markerballs.	67
2.29	Histogram of the measures of Δ over the database, indicating the average difference between the transformations.	67
2.30	Comparison of the linear and the non-linear registration on two sequences where the difference is of category B. (a) Sequence 1 processed with the linear registration, (b) Sequence 1 processed with the <i>non</i> linear registration, (c) Sequence 2 processed with the linear registration, (d) Sequence 2 processed with the <i>non</i> linear registration.	68
2.31	Histogram of StentViz image score - StentOp Image score; Hypothesis H_0 is represented by the blue dot. 95% confidence interval for the difference of scores is represented in red (rated by: Left reader 1, center reader 2 and right reader 3).	71
2.32	Two examples ((a) and (b)) and the images produced by StentViz ((c) and (d)) and StentOp ((e) and (f)). Note that with StentViz the details of the stent are more visible.	72
2.33	DSE image: Observe the plaque indicated by the arrow that prevents the proper deployment of the stent.	76
2.34	Use of DSE to optimize stent overlapping.	77
2.35	Use of DSE for a bifurcation treatment procedure. Top: bifurcation before treatment presenting significant stenoses. Center row: a first stent is deployed / a second stent is positioned next to it / both stent expanding (note the overlapping) / stent shape is further modified by balloon inflation in the side branch. Bottom: Final angiographic result of the treated bifurcation.	78
2.36	Use of DSE to detect stent fracture. Left: a coronary aneurysm is indicated by the arrow. Center: DSE enables to see that the aneurysm is linked with the stent rupture indicated by the arrow. Right: after treatment.	79

2.37	Illustration of the variability of the appearance due to strut superimpositions patterns. (a) and (b) depict a same stent under two viewpoints differing from few degrees.	79
2.38	Illustration of imperfect registration in some cases (see text for details). (a) Although the landmarks are well registered and sharp, the stent struts do not appear and the upper part of the stent is blurry. (b) Case of two stents at a bifurcation. Only the marker-balls and the guide-wire supporting them are used as registration landmarks. Therefore, the stent in the side branch is not registered as accurately as the one in the main branch and is blurrier.	80
3.1	(a) input image. (b) output image of StentViz "guide-wire subtracted", top regular DSE image, bottom guide-wire subtracted.	81
3.2	Scheme illustrating the difficulties in stent visualization introduced by the presence of the guide-wire. Top: the guide-wire hides the structure of the stent struts; Middle: the guide-wire hides partially the border of the stent; Bottom: ideal case for stent visualization where the guide-wire is not displayed (images are adapted from a Cypher stent photo [tim05]).	83
3.3	Stent visualization difficulties in clinical DSE images. Left: the guide-wire hides the structure of the stent struts; Right: the guide-wire hides partially the border of the stent.	83
3.4	Illustration of the continuous (volumetric) Xray image formation model (left) and the discrete planar layer model (right).	84
3.5	Planar layer modeling of a DSE image: 3 separate layers, namely the anatomical background, the guide-wire and the stent, are combined in the image formation mechanism.	87
3.6	Background subtraction in DSE images. Left: motion compensated integration of the images. Right same as left performed on background subtracted images.	88
3.7	Anatomy of a coronary guide-wire illustrating its main components: the core, the cover and the tip (Figure reproduced from [ENS ⁺ 10]).	89
3.8	Theoretical attenuation profile of a uniform cylindrical guide-wire imaged with a perfect system ($MTF = \text{dirac}$). Abscissa: distance to the centerline in pixels (the radius of the guide-wire is here 1,5 pixels). Ordinates: attenuation factor (no dimension).	89
3.9	From left to right: DSE image, guide-wire subtracted image, estimated guide-wire layer, curve of the intensities of the pixels on the guide-wire in the DSE image. Observe on the curve all the variations induced by the stent signal as a consequence of the transparency of X-ray images.	91
3.10	From left to right: input image, estimated guide-wire layer, guide-wire subtracted image.	91
3.11	From left to right: regular DSE image, guide-wire subtracted DSE image, both downscanned to fit the display, both cropped on the area of interest to fit the display.	91
3.12	Simultaneous display of two DSE image with/without guide-wire subtraction in the case of a diagonal stent. The area of interest is represented by the union of the three squares in the upper part.	93

3.13	Attenuation profiles of 10 guide-wires in clinical cases estimated with the average (left) and the median (right). Abscissa: distance to the guide-wire centerline in pixels. Ordinate: attenuation factor (no dimension). The average profile with both methods is plotted with a dashed line. . . .	94
3.14	Guide-wire subtraction results: left average, right median. Observe on the image on the left the white line (artifact) at the center of the stent. .	95
3.15	(a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.	97
3.16	(a) input image. (b) guide-wire subtracted image. The guide-wire subtraction is satisfactory but not perfect. We can observe in the upper section of the stent a residual guide-wire signal. Our hypothesis to explain it is that the guide-wire centerline estimation is less reliable in this area where the stent border and the guide-wire are very close one to another. .	97
3.17	(a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction. The stent-signal is rather weak compared to the intensity of the guide-wire, therefore subtraction improves one's ability to assess the general shape of the stent.	97
3.18	(a) input image. (b) guide-wire subtracted image. The subtraction is not perfect in the lowest part of the stent where one can notice an unusually white signal. However, on the upper part of the stent, the subtraction drives the attention of the reader on a dark area. Its presence is also visible (but far less conspicuous) in the non-subtracted image. This area is likely to be a large calcification that can have impacted stent deployment in a plane perpendicular to the one of the image. Such a finding can encourage the cardiologist to repeat a StentViz acquisition with a different viewing angle.	97
3.19	(a) input image. (b) guide-wire subtracted image. The subtraction proved to be efficient in this stent in stent situation (two stents are superimposed in the lowest part). StentViz enables to understand the relative positioning of the two stents.	98
3.20	(a) input image. (b) guide-wire subtracted image. Perfect subtraction example.	98
3.21	(a) input image. (b) guide-wire subtracted image. Excellent guide-wire subtraction.	98
3.22	(a) input image. (b) guide-wire subtracted image. Excellent guide-wire subtraction.	98
3.23	(a) input image. (b) guide-wire subtracted image. Interesting case where the guide-wire is constantly at the same distance to the stent border. Thanks to the median estimator in the estimation of the attenuation profile the processing is robust to such situations.	99
3.24	(a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.	99
3.25	(a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.	99
3.26	(a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.	99
3.27	This case is of particular clinical interest. The guide-wire subtraction (on the right) reveals the border of the stent that is not straight, as indicated by the arrow. This finding was not obvious in the regular DSE image (left). .	100

3.28	(From left to right: Regular contrast setting (no saturation); Aggressive contrast setting allows to improve the display of the stent but saturate the guide-wire; With guide-wire subtraction this contrast setting does not create saturation in the stent area.	100
4.1	3D coronary stent reconstruction at a bifurcation.	102
4.2	Illustration of a locally under-deployed stent in a pig artery. In image (a) the under-deployment circled in blue is visible. In image (b), under a different viewing angle the defect is hardly noticeable (images from [Per08]).	103
4.3	(a) A frame of the input sequence. (b) Visualization of the stent reconstruction with the method of Perrenot [Per08]. (c) Same stent reconstructed with a CT scan. Images extracted from [Per08].	104
4.4	Five images extracted from a rotational acquisition. Note the changes in balloon appearance and marker position. In the second image, one of the markers is overlapping the tip of the catheter, making it hardly visible. In the third image sternal wires appear, and in the fourth the markers overlap the guide-wire tip.	105
4.5	General overview of the marker tracking algorithm.	107
4.6	Merging two potentially overlapping tracks A and B is reduced to merging non overlapping ones by assigning all the common indexes to either one or the other.	110
4.7	Two non overlapping tracks C and D are represented. The merging point couples $(\xi(C, D)$ in the text) are represented with a wider black border.	110
4.8	Plot of the length of the detected marker segment versus the frame index. False positives are circled in green.	112
4.9	(a) et (b): Frames 158 and 160 of a clinical sequence. The actual markers have been detected at the locations indicated in green. However, in the intermediate frame (number 159) they have not been detected. (c) The set of detected pairs in frame 159. The actual marker pair does not belong to them. (d) The ROI for pairing and the new set of pairs in frame 159, based on the detected markers in the previous and next frames. (e) Extract from (c) where we can observe that the actual markers did not belong to the set of detected pairs.	114
4.10	Top: Example of wrong marker labeling. Bottom: A typical case where marker labeling is a difficult case because the markers invert their relative positions along the sequence.	115
4.11	3D reconstruction of the center point from 3 detections.	116
4.12	Plane associated with a marker couple i	116
4.13	Projection of the rough 3D model on the image, translation of the projected segment and labeling of the detected marker couple (in green).	118
4.14	3 images of a spin sequence. Top row: original image. Central row: segmented guide-wire. Bottom row: subtracted guide-wire.	119

- 4.15 A situation where affine registration is not optimal. Between images (a) and (b) we assume that the marker segment has undergone a rotation out of the image plane, around the (Ox) axis. We illustrate in red the marker segment of image (a). Note the change in the marker segment length between figures (a) and (b). The change in distance between the markers in images (a) and (b) is compensated by a global scaling with the similarity transform that modifies the (nearly unchanged) diameter of the stent (image (c)). In image (d) a stretching along the axis of the markers but not perpendicularly keeps unchanged the stent diameter while registering the markers. In these two last images the registered markers are illustrated in green. 121
- 4.16 Motion compensated images. Left and center: Two registered images in a phase where the gantry is static. Right: Registration of the manually segmented stents. Top row state of the art similarity transform, bottom row linear diameter preservation registration. Observe that the registered stents match better with the proposed registration than with the classical affine one (as indicated by the arrows). 122
- 4.17 Illustration of the non-linear registration and of the guide-wire subtraction on three images taken at different instants during a rotational acquisition. We represent the segmented guide-wire in green, and the projected 3D guide-wire model in red. Top row: original images. Bottom row: registered and guide-wire subtracted images. Left: We can observe the discontinuity of the guide-wire based registration far from the segmented guide-segment (notice the “broken guide-wire effect”). Middle: the projected and detected guide-wires occasionally match very closely, typically once per heart beat, as in this example. Right: notice the difference in curvature between the projected and detected guide-wires. 124
- 4.18 Automated marker detection and labeling on sequence 3. 126
- 4.19 Left: 3D stent model reconstructed from sequence 3. Right: 3D stent of sequence 6 (bifurcation case) 127
- 4.20 3D stent reconstructions from left to right: sequence 2, 3a and 1a. 127
- 4.21 Left longitudinal slice / Right axial slice of the reconstructed stent (image from [Per08]). 128
- 4.22 3D stent slice: Left classical processing, right guide-wire subtracted between the markers. The stent border hidden by the guide-wire on the left is visible on the right, as indicated by the white arrows. The guide-wire subtraction process attenuates the guide-wire signal, but in the 3D volume a mark still persists. 128
- 4.23 Left: state of the art. The transform to register the markers is a similarity and the guide-wire is not subtracted; center: radius preserving affine registration and guide-wire subtraction; right: non-linear registration and guide-wire subtraction. The three rows illustrate different viewing angles of the same stent. Observe on top right that the strut pattern appears in the upper side of the stent. This pattern is not visible on the left image. On the bottom right the general shape of the stent can be easily assessed (this it is not the case on the left). 129
- 4.24 Comparison of the similarity transform (right) and of the non-linear one (left), in the slices of a reconstructed stent volume (same example as Fig. 4.23. Guide-wire has been subtracted in both cases. 130

4.25	Stability of the estimation of the parameters of the 2D registration versus the marker segment length. Left: inverse of the marker segment length (in pixels) vs the frame index. Right: estimated rotation angle in radians vs the frame index.	132
4.26	Volume rendering of a same stent, with (left) and without (right) balloon inflation. Observe that the inflation of the balloon is (i) beneficial to the quality of the enhanced stent and that (ii) prevents the guide-wire from overlapping with the stent border. This illustration case is the one where the impact of balloon inflation is the most striking in our database.	132
5.1	Confusion matrix	139
5.2	ROC space (image from [Ind09]). Point B illustrate a random choice classifier. C' is a better classifier than A. C is the negation of C'.	141
5.3	ROC curve illustration (image from [Ind09]).	141
5.4	Illustration of the computation of TP, FP, TN and FN. Top row: original image, ground truth and dilated ground truth. Bottom row, ROI, output of our line enhancement technique and threshold and skeletonization of this output.	145
5.5	X-ray fluoroscopy image. A guide-wire is illustrated. Note that it is a long smooth curve, with low contrast to noise ratio.	148
5.6	Illustration of the definition of the paths by Vincent. The center pixel, where the path originates, is p . l is the length of the path, θ_i its orientation, and the angular range $\delta\theta$ the straightness. The search cone in which paths are constrained to be located is shaded on the figure.	150
5.7	Illustration of the recursivity of the cone constraint. A path originating from a pixel p shall stay in the cone placed in any of its pixels p' . The path of (a) verifies this property illustrated in (b). The path of (c) does not verify this property (see (d)) and will thus not be considered by the technique.	150
5.8	Triangular search cones defined by Vincent. These four cones of straightness $\pi/2$ cover the whole angular range.	150
5.9	Left a curvilinear structure in a very noisy environment. Right, minimal path computed between the top and the bottom of the image (illustration from [Vin98]). It is extremely tortuous, and does not follow the smooth curves that are naturally present in medical images.	151
5.10	500 paths originating at random locations are illustrated. We can observe how they all tend to converge to the linear structures of the image and especially to the guide-wire.	154
5.11	Extract of the top image. Observe the density of paths on the guide-wire.	154
5.12	Here we built a test image with a perfect circle. We illustrated again the content of the polygonal path image at 500 random locations. The behavior is similar to the one observed on the guide-wire image.	154
5.13	We illustrated 5000 paths originating at random locations. We can observe on the line that is similar to the medial axis of the curvilinear structure for our problem, that the paths are attracted by different equidistant parts of the guide-wire. We approximately represented this line in black on the figure.	155
5.14	All the paths intersecting at one given point. This particular point belongs to the guide-wire.	155

5.15	Illustration of the notations for the computation of \mathfrak{P} . Let the center pixel be of coordinates (i, j) . In the image I_{θ_k} we store the sum of the pixels along line segments of orientation θ_k . Such a line segment is illustrated on the figure with plain lines. Translation parameters (t_{x_k}, t_{y_k}) are also stored. A search cone is defined by a start angle <i>cone.start</i> and an end angle <i>cone.end</i> . All the possible locations of the end of the segments in a given search cone are illustrated by the arc of circle in dashed line. These are the locations of the pixels that will be evaluated to extend the path in the successive iterations.	157
5.16	At the second iteration of the algorithm, each pixel p_i contains the minimal path originating from it of length l . In order to compute the minimal path of 2 segments (length $2l$) originating from pixel p we consider each pixel p_i and note that the minimal path originating from p and passing through p_i is simply the already computed path originating from p_i to which we add the line segment $p_i p$	157
5.17	We illustrated 3 typical paths of 2 segments that have been considered by the algorithm.	158
5.18	On a synthetic image representing a circle we have computed \mathfrak{P} and associated to each pixel the cost of the path originating from it. We can observe that the image is complex and that linear structures tend to extend to neighbor pixels.	159
5.19	We have computed \mathfrak{P} on the clinical from Fig. 5.5 and associated to each pixel the cost of the path originating from it. We can observe that the image is complex and that linear structures tend to extend to neighbor pixels.	159
5.20	Path voting on the image illustrated on Fig. 5.5. See section 5.5.3 for details on the pre-processing.	160
5.21	A more aggressive contrast setting (saturation) of fig 5.20 shows lower values of the image. We see some path that are attracted by the guide-wire. Especially close to the tip.	160
5.22	Even more aggressive contrast setting of image fig 5.20 enables to see its complex structure.	160
5.23	Similar image as fig 5.20 except that constrained path smoothness, rejecting any path of tortuosity lower than 0.6	161
5.24	A more aggressive contrast setting (saturation) of fig 5.23 shows lower values of the image. We see some path that are attracted by the guide-wire. Especially close to the tip, however there are less such path than in 5.21 thanks to the tortuosity constraint.	161
5.25	Even more aggressive contrast setting of image fig 5.20 enables to see its complex structure. Compared to fig 5.22 we see that the most tortuous path have been removed.	161
5.26	$\vartheta(\mathfrak{P})$ with 3 pixel long polygonal segments (top) and 11 pixel long segments (bottom) for a similar total path length. In the left column (nominal windowing) we observe that the curvilinear structure appear tortuous with short path segments and smooth with longer segments (as in the original image).	162

5.27	We illustrate here the result of path pruning with $x = 50\%$. We can observe that among the small number of paths selected, the ones illustrating the guide-wire are present. Further path selection, based on their cost for instance can be performed to better isolate the guide-wire.	163
5.28	We illustrate here the result of path pruning with $x = 50\%$ and tortuosity constraints (top 0.6, bottom 0.75). We can observe how the conjunction of pruning and constrained path smoothness help segment the guide-wire.	164
5.29	Extract of $\vartheta(\mathfrak{P})$ obtained from the X-ray image of a guide-wire. The arrow resulting from direction estimation give the direction of the closest curvilinear structure. Note that on the structure itself the direction of the arrow is not relevant (the linear structure can be found in both opposite directions), but the orientation is accurate.	165
5.30	Counter example to the triangular inequality when path curvature is bounded. The “distance” between two points is defined as the minimal cost of all the paths joining two points, under the constraint that the maximum curvature of the paths shall not exceed a given bound. We build a particular potential cost image for this counter example, where: all the image has cost $+\infty$ except the three paths that are drawn. Let L_{AB} be the length of the path from A to B. We assign to each pixel on this path the cost $3/L_{AB}$. Consequently the cost of the whole path is 3. Similarly we assign to the pixels on the path from A to C the cost $1/L_{AC}$ and to those on the path from C to B the cost $1/L_{CB}$. The cost of each of these two paths is thus 1. Then, we note that the only finite cost path from A to B, is the one that is drawn (see Fig. (c)). Indeed, the only other finite cost path is the one going through C, but it does not verify the curvature bound constraint. Therefore $d(A, B) = 3$. Similarly $d(A, C) = d(C, B) = 1$. Finally: $d(A, B) > d(A, C) + d(C, B)$	167
5.31	The cone distance.	167
5.32	Counter example for the triangular inequality on λ . We consider the 4 cones used by Vincent [Vin98] that we group in two pairs of opposite cones, illustrated on (a) and (b). We build a particular potential cost image (b) where: all the image has cost ∞ except the three paths that are drawn. Let L_{AB} be the length of the path from A to B, we assign to each pixel on this path the cost $3/L_{AB}$. Consequently the whole path cost is 3. Similarly we assign to the pixels on the path from A to C the cost $1/L_{AC}$ and to those on the path from C to B the cost $1/L_{CB}$. The cost of each of these two paths is thus 1. Then we note that the only finite cost path from A to B is the one that is drawn that is included in cone (a) (see Fig. (c)). Therefore $\lambda(A, B) = 3$. Similarly $\lambda(A, C) = 1$. The only finite cost path from C to B is the one that is drawn, and it is in the cone (b) (see Fig. (d)). Consequently $\lambda(C, B) = 1$. Finally: $\lambda(A, B) > \lambda(A, C) + \lambda(C, B)$	170
5.33	Illustrations of $\mathcal{R}_{A,B}$, for several configuration of A and B with the set of four non-overlapping cones presented in Fig 5.8. In this case $\mathcal{R}_{A,B}$ is the inter-visibility parallelogram.	170
5.34	Images extracted from [BC09]	171
5.35	Application of the PSI to clinical examples.	174
5.36	Application of the PSI to clinical examples.	175
5.37	Application of the PSI to clinical examples.	176

5.38	Application of the PSI to clinical examples.	177
5.39	Illustration of the performance of the line enhancement techniques (measured by the partial AUC) on each sequence of the database. Top: FPR is computed on the whole image. Bottom: FPR is computed only in a band around the guide-wire.	178
5.40	Average performance of four line enhancement techniques over our database of clinical cases. Top: FPR is computed on the whole image. Bottom: FPR is computed only in a band around the guide-wire.	178
5.41	MDR (blue), FDR (red) and MDR+FDR (green) as a function of the threshold in three different clinical cases. As expected, MDR increases, FDR decreases and their sum has a clear minimum. In this experiment we chose 500 thresholds regularly sampled between the max and the min of the image.	179
5.42	Top: input image. Middle row: Frangi (left) and RLT (right). Bottom row: $\vartheta(\mathfrak{P})$ (left) and $\vartheta_{\tau_{\min}}(\mathfrak{P})$ (right). Observe that traditional methods (middle row) enhance a large amount of non relevant structures in the background. Moreover, they fail to enhance some parts of the guide-wire. These problems are not present with $\vartheta_{\tau_{\min}}(\mathfrak{P})$	180
5.43	Similar experiment as Fig. 5.42 on another clinical case. This one is the most challenging case yielding the poorest results due to the presence of the ribs that are very visible under this viewing angle and of sternal wires (first sequence in Fig. 5.39). Same observations apply.	181
5.44	Similar experiment as Fig. 5.42 on another clinical case. Same observations apply.	182
5.45	Similar experiment as Fig. 5.42 on another clinical case. Same observations apply.	183
5.46	Similar experiment as Fig. 5.42 on another clinical case. Same observations apply.	184

List of Tables

1.1	Some key figures about PCI (typical orders of magnitude).	8
1.2	Steps of a stenting procedure. First column is the index of the step, second one its description, third one the acquisition mode (F stands for fluoro and C for cine), and the last one the clinical challenges. Steps 3 to 6 can be omitted in case of direct stenting.	12
2.1	Repartition of the images produced with the different estimators according to relative image quality scale (percentages).	42
2.2	Stents in the image quality study: stent model, manufacturer, number of occurrences, input image quality, image quality after the application of DSE and improvement in image quality.	62
2.3	Influence of the transform	68
2.4	Stents types in the study	69
2.5	Stents dimensions in the study	70
2.6	Image quality ratings	70
2.7	Comparative ratings	71
4.1	Result of the detection and tracking algorithm designed for DSE.	105
4.2	Results of the algorithm vs ground truth for each different stage.	125
5.1	Comparison of the line enhancement techniques. SOD stands for Second Order Derivative, RLT for Rotated Line Template and SF for Steerable Filter. These are generic categories of filters that include a large number of variations (see references in the table for more details).	146

Glossary

angina pectoris commonly known as angina, is severe chest pain due to ischemia (a lack of blood, hence a lack of oxygen supply) of the heart muscle, generally due to obstruction or spasm of the coronary arteries (the heart's blood vessels). Coronary artery disease, the main cause of angina, is due to atherosclerosis of the cardiac arteries. The term derives from the Latin *angina* ("infection of the throat") from the Greek *ankhōnē* ("strangling"), and the Latin *pectus* ("chest"), and can therefore be translated as "a strangling feeling in the chest". Worsening ("crescendo") angina attacks, sudden-onset angina at rest, and angina lasting more than 15 minutes are symptoms of unstable angina (usually grouped with similar conditions as the acute coronary syndrome) (from <http://encyclopedia.thefreedictionary.com/Angina+pectoris>). 1

angiogenesis The process of developing new blood vessels (from <http://www.medterms.com/script/main/art.asp?articlekey=6597>). 27

angiogram or angiography is a medical imaging technique used to visualize the inside, or lumen, of blood vessels and organs of the body, with particular interest in the arteries, veins and the heart chambers. This is traditionally done by injecting a radio-opaque contrast agent into the blood vessel and imaging using X-ray based techniques such as fluoroscopy. The word itself comes from the Greek words *angeion*, "vessel", and *graphein*, "to write or record". The film or image of the blood vessels is called an angiograph, or more commonly, an angiogram (from <http://en.wikipedia.org/wiki/Angiography>). 4, 9, 15

angiographic catheter also called injection catheter. A catheter through which a contrast medium is injected for visualization of the vascular system of an organ (from <http://medical-dictionary.thefreedictionary.com/Angioplasty,+balloon>). 9

atherosclerosis (also known as arteriosclerotic vascular disease or ASVD) is a condition in which an artery wall thickens as the result of a build-up of fatty materials such as cholesterol. It is a syndrome affecting arterial blood vessels, a chronic

inflammatory response in the walls of arteries, in large part due to the accumulation of macrophage white blood cells and promoted by low-density lipoproteins (plasma proteins that carry cholesterol and triglycerides) without adequate removal of fats and cholesterol from the macrophages by functional high density lipoproteins (HDL). It is commonly referred to as a hardening or furring of the arteries. It is caused by the formation of multiple plaques within the arteries (from <http://encyclopedia.thefreedictionary.com/Atherosclerosis>). 3

balloon catheter a catheter whose tip has an inflatable balloon that holds the catheter in place or can dilate the lumen of a vessel, such as in angioplastic procedures (from <http://medical-dictionary.thefreedictionary.com/Angioplasty,+balloon>). 6

balloon markerballs the balloon catheter has metallic markers (either at the center or on either side of the balloon). This helps the cardiologist know the location of the otherwise "invisible" balloon (from <http://www.heartsite.com/html/ptca.html>). 16

bare metal stent a vascular stent without a coating (as used in drug-eluting stents). It is a mesh-like tube of thin wire. The first stents licenced for use in cardiac arteries were bare metal - often 316L stainless steel. More recent ('2nd generation') stents use cobalt chromium alloy (from http://en.wikipedia.org/wiki/Bare-metal_stent). 9

catheterization lab or cathlab, the room where the catheterization procedure takes place. 9

chronic total occlusion (CTO) Although it is often impossible to know exactly when the occlusion occurred, a CTO is arbitrarily defined as a > 3 month old, total obstruction of a coronary artery (from http://heart.bmj.com/content/91/suppl_3/iii42.extract). 27

coronary angiography also called coronarography. The most accurate method (the gold standard) for evaluating and defining coronary artery disease (CAD). Coronary angiography is used to identify the exact location and severity of CAD. 8

coronary stent a tube placed in the coronary arteries that supply the heart, to keep the arteries open in the treatment of coronary heart disease. It is used in a procedure called percutaneous coronary intervention (PCI). Stents reduce chest pain and have been shown to improve survivability in the event of an acute myocardial infarction. Similar stents and procedures are used in non-coronary vessels e.g. in the legs in peripheral artery disease. 1

drug eluting stent (DES) a peripheral or coronary stent (a scaffold) placed into narrowed, diseased peripheral or coronary arteries that slowly releases a drug to block cell proliferation. This prevents fibrosis that, together with clots (thrombus), could otherwise block the stented artery, a process called restenosis. The stent is usually placed within the peripheral or coronary artery by an Interventional Cardiologist or Interventional Radiologist during an angioplasty procedure (from http://en.wikipedia.org/wiki/Drug-eluting_stent). 8

guide-wire a thin, usually flexible wire that can be inserted into a confined or tortuous space to act as a guide for subsequent insertion of a stiffer or bulkier instrument (from <http://medical-dictionary.thefreedictionary.com/guidewire>). 11

interventional cardiology is a branch of the medical specialty of cardiology that deals specifically with the catheter based treatment of structural heart diseases. This most commonly involves the insertion of a sheath into the femoral artery and cannulating the heart under X-ray visualization. The interventional cardiology procedure of primary angioplasty is now the gold standard of care for an acute myocardial infarction. It involves the deployment of stents and balloons, leaving no scars, which has given it the name "pin-hole surgery" (as opposed to "key-hole surgery") (from <http://encyclopedia.thefreedictionary.com/Interventional+cardiology>). 1

ischemia (from Greek, ischaima; isch- root denoting a restriction or thinning or to make or grow thin/lean, haema blood) is a restriction in blood supply, generally due to factors in the blood vessels, with resultant damage or dysfunction of tissue. It may also be spelled ischaemia. It also means local anemia in a given part of a body sometimes resulting from congestion (such as vasoconstriction, thrombosis or embolism (from <http://encyclopedia.thefreedictionary.com/Ischemia>). 4

myocardial infarction (MI) or acute myocardial infarction (AMI), commonly known as a heart attack, is the interruption of blood supply to a part of the heart, causing heart cells to die. This is most commonly due to occlusion (blockage) of a coronary artery following the rupture of a vulnerable atherosclerotic plaque, which is an unstable collection of lipids (fatty acids) and white blood cells (especially macrophages) in the wall of an artery. The resulting ischemia (restriction in blood supply) and oxygen shortage, if left untreated for a sufficient period of time, can cause damage or death (infarction) of heart muscle tissue (myocardium). 1, 9

PCI Percutaneous Coronary Intervention. 6

restenosis literally means the reoccurrence of stenosis, a narrowing of a blood vessel, leading to restricted blood flow. Restenosis usually pertains to an artery or other large blood vessel that has become narrowed, received treatment to clear the blockage and subsequently become renarrowed. This is usually restenosis of an artery, or other blood vessel, or possibly a vessel within an organ. This term is common in vascular surgery, cardiac surgery, and angioplasty, all branches of medicine that frequently treat narrowing of blood vessels. 6

stem cell One of the human body's master cells, with the ability to grow into any one of the body's more than 200 cell types. All stem cells are unspecialized (undifferentiated) cells that are characteristically of the same family type (lineage). They retain the ability to divide throughout life and give rise to cells that can become highly specialized and take the place of cells that die or are lost (from <http://www.medterms.com/script/main/art.asp?articlekey=10597>). 27

stenosis (plural: stenoses) is an abnormal narrowing in a blood vessel or other tubular organ or structure (from <http://en.wikipedia.org/wiki/Stenosis>). 1

stent In the technical vocabulary of medicine, a stent is an artificial 'tube' inserted into a natural passage/conduit in the body to prevent, or counteract, a disease-induced, localized flow constriction. The term may also refer to a tube used to temporarily hold such a natural conduit open to allow access for surgery (from <http://en.wikipedia.org/wiki/Stent>). 1

StentViz a stent enhancement visualization software (DSE) commercialized by GE. 17

thrombosis Thrombosis is the formation of a blood clot inside a blood vessel, obstructing the flow of blood through the circulatory system. When a blood vessel is injured, the body uses platelets and fibrin to form a blood clot to prevent blood loss. Alternatively, even when a blood vessel is not injured, blood clots may form in the body if the proper conditions present themselves. If the clotting is too severe and the clot breaks free, the traveling clot is now known as an embolus (from <http://encyclopedia.thefreedictionary.com/Thrombosis>). 6

thrombus A clot consisting of fibrin, platelets, red blood cells, and white blood cells that forms in a blood vessel or in a chamber of the heart and can obstruct blood flow. The rupture of atherosclerotic plaques can cause arterial thrombosis (the formation of thrombi), while tissue injury, decreased movement, oral contraceptives, prosthetic heart valves, and various metabolic disorders increase the risk for venous thrombosis. A thrombus in a coronary artery can cause a heart attack (from <http://www.thefreedictionary.com/thrombus>). 4

Xray X-radiation (composed of X-rays) is a form of electromagnetic radiation. X-rays have a wavelength in the range of 0.01 to 10 nanometers, corresponding to frequencies in the range 30 petahertz to 30 exahertz) and energies in the range 120 eV to 120 keV. They are shorter in wavelength than UV rays and longer than gamma rays. In many languages, X-radiation is called Rontgen radiation, after Wilhelm Conrad Rontgen, who is generally credited as its discoverer, and who had named it X-radiation to signify an unknown type of radiation (from <http://encyclopedia.thefreedictionary.com/X-ray>). 1

Xray fluoroscopy real-time imaging using an x-ray source that projects through the patient onto a fluorescent screen or image intensifier. Image-intensified fluoroscopy has replaced conventional fluoroscopy in current practice (from <http://medical-dictionary.thefreedictionary.com/x-ray+fluoroscopy>). 1, 9

Chapter 1

Introduction

This PhD thesis in the field of image processing falls into the context of curing coronary heart disease, the leading cause of death for both men and women worldwide. One of the most common treatment options is a minimally invasive procedure monitored under [Xray fluoroscopy](#). It consists in the deployment of a [stent](#) (a fine scaffolding of wire), in the clotted artery to re-open it and restore the blood flow. The proper deployment of the stent and its apposition to the vessel wall are key to the success of the intervention. The purpose of this thesis is to study ways of imaging [coronary stents](#) with Xray fluoroscopy, in order to assess qualitatively and quantitatively their shape and deployment. In this introductory chapter, we first present coronary heart disease and its most common treatment, stenting, in [Section 1.1](#). Then we expose the clinical facts motivating the need to visualize stents during the course of an intervention and we summarize the ultimate needs of the clinicians regarding stent visualization ([Section 1.2](#)). [Section 1.3](#) covers the limitations of the existing stent visualization techniques and grounds the need for new ones. The technical context of this thesis is described in [Section 1.4](#). [Section 1.5](#) presents the overall structure of the thesis, describing how each chapter answers to the identified needs and limitations. Our main technical contributions are summarized in [Section 1.6](#). Finally, [Section 1.7](#) covers some perspectives for the near future of the cathlab.

1.1 Pathology and treatment

This section aims at giving to the reader a brief overview of the clinical context and introducing the terms that are commonly met in the field of [interventional cardiology](#). We introduce the cardiovascular system, then we explain the pathology of atherosclerotic [stenosis](#) and its potential consequences: [angina pectoris](#) and [myocardial infarction](#). Finally we present the associated treatment options.

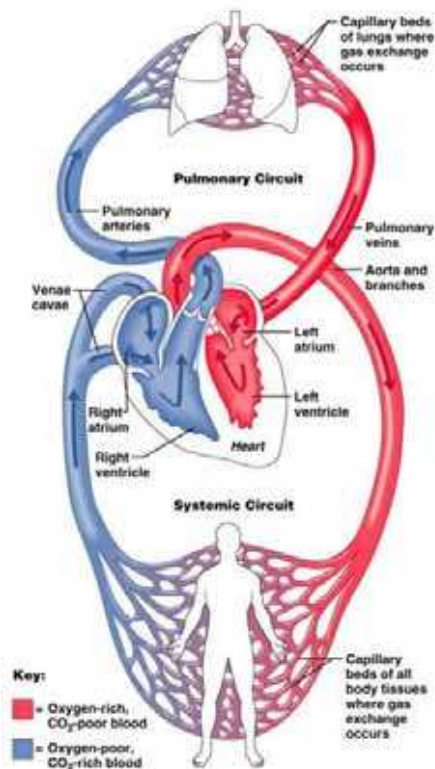


FIGURE 1.1: Simplified view of the cardiovascular system (Copyright © Addison Wesley Longman, Inc publishing as Benjamin Cummings).

1.1.1 The cardiovascular system

The cardiovascular system is in charge of the circulation of blood through the entire body providing oxygen and nutrient supply to the organs. It is constituted of the heart which similarly to a pump is responsible for the circulation of the blood, and of the arteries, veins and capillaries through which blood is carried. The blood circulates through the body inside two circuits that interconnect at the level of the heart (Fig. 1.1):

- the *pulmonary circulation* which carries oxygen-depleted blood away from the heart, to the lungs, and returns oxygenated blood back to the heart at the entry of the systemic circulation.
- the *systemic circulation* which carries oxygenated blood away from the heart to the body, and returns deoxygenated blood back to the heart at the entrance of the pulmonary circulation.

Just like every other organ, the heart has its own vascularisation system formed of veins and arteries. The coronaries arteries, whose name comes from their location as a crown around the heart, are irrigating, thus feeding the heart muscle called the myocardium. The coronary arteries are formed of two distinct coronary trees, the right one and the left one that both originate from the beginning of the aortic root immediately

above the aortic valve. The left coronary artery that is generally more developed than the right one, divides into two main branches: the left anterior descending (LAD) and the left circumflex artery (LCX). The Figure 1.2 depicts the coronary arteries.

The coronary arteries are classified as *end circulation*, since they represent the only source of blood supply to the myocardium: there is very little redundant blood supply, which is why blockage of these vessels can be so critical [Wik11a]. We will see in the next section that these relatively narrow vessels are commonly affected by a pathology named [atherosclerosis](#) and can see their diameter drastically reduced, causing angina or heart attack.




FIGURE 1.2: Overview of the coronary arteries [LtOH10].

Each beating of the heart produces the following chain of events, called heart cycle:

- **atrial systole**, the atria contract and eject blood into the ventricles.
- **ventricular systole** contraction of the ventricles injecting blood into the circulatory system. The atria, now relaxed are being filled with blood.
- **diastole**, relaxation of all parts of the heart allowing ventricles to be filled.

1.1.2 The atheromatous stenosis and the myocardial infarction

Coronary heart disease (CHD) and its consequences are one of the major world public health issues, as we can understand from the following few figures on its epidemiology. According to the World Health Organization [Org11] 7.2 million people died from coronary artery disease in 2004 representing 12% of all deaths. Moreover it is the leading cause of death in developed countries. For instance, statistics for 2007 in the United States [RGLJA11] show that it is the single leading cause of death, representing one

sixth of the deaths. Approximately 16,000,000 Americans over the age of 20 have CHD (7% of the population). The two main consequences of CHD are heart attack that affects 1,225,000 people per year and angina that affects 10,200,000 living people both in the United States.

We are particularly interested here in the main cause of angina and heart attack: the atheromatous stenosis. This disease is also called atherosclerosis. It is the gradual buildup of cholesterol and fibrous tissue in plaques in the wall of coronary arteries typically over decades (Fig. 1.3). Blood stream column irregularities visible on [angiography](#) (Fig. 1.4), an Xray technique that enables vessel visualization, reflect artery lumen narrowing as a result of decades of advancing atherosclerosis. When the narrowing is severe enough, it can have consequences on the blood supply of the depending heart territory. It typically results in a severe chest pain due to [ischemia](#) (a lack of blood, hence a lack of oxygen supply) of the heart muscle. It is called angina pectoris (commonly known as angina). It often appears under stress, when the demand of oxygen from the heart muscle is higher, and stops at rest. From a quantitative standpoint, a stenosis is considered significant if the narrowing obstructs more than 70% of the diameter of the artery (Fig. 1.5) and non significant below 50% [BLA02]. For intermediate cases, additional information is necessary to quantify the impact of the stenosis.

At a more severe stage, plaques can become unstable, rupture, and additionally promote a [thrombus](#) (blood clot) that occludes the artery; this can occur in minutes [Wik11c]. When a severe enough plaque rupture occurs in the coronary vasculature, it leads to myocardial infarction (necrosis of downstream myocardium), known as heart attack, that can cause death. The non irrigated cardiac cells can eventually die, typically after 4 hours, causing non-reversible damage since these cells do not regenerate themselves. A collagen scar forms in their place. The injured tissue will not contract similarly to normal tissue, nor will it conduct properly electrical impulse, putting the patient at risk for potentially life threatening complications. As a rule of thumb, the longer the infarct, the greater the damage. It is summarized in the famous adage “**time is muscle**” [GDLA04, Ant08].

1.1.3 Treatment options

The main treatment options share some similarities for the angina and the acute myocardial infarction. There are four axes in these therapies: drug administration, minimally invasive procedures, bypass, and risk factor management.

1.1.4 Drug administration

The drugs commonly involved in the treatment of the angina are inspired by the following strategy:




FIGURE 1.3: Progression of the atheromatous stenosis over decades, potentially rupturing putting the patient at risk of major adverse events (image from [Chi06]).




FIGURE 1.4: Visualization of a diseased artery under Xray fluoroscopy (angiography) with an Innova 2100 Xray system (GE)

- Reduce the needs in oxygen of the heart, slow its rate and decrease arterial pressure (Beta blockers and calcium channel blockers).
- Help dilate the coronary arteries to improve the blood supply (vasodilators drugs like nitrates).
- Prevent the formation of blood clots (Antiplatelet/anticoagulant drugs such as aspirin).




FIGURE 1.5: Quantitative coronary analysis allows the clinician to estimate the percentage of lumen reduction induced by the presence of the stenosis (Quantitative Analysis Software developed by Pie Medical Imaging BV ©; pictures courtesy by Pie Medical Imaging BV. Reproduced from [esa11]).

If the patient is presumed to have an occlusive [thrombosis](#) (typically in case of myocardial infarction) a specific drug is envisaged called thrombolysis that dissolves the thrombus.

1.1.5 Angioplasty

The angioplasty, or Percutaneous Coronary Intervention ([PCI](#)) is a minimally invasive procedure that consists in introducing a [balloon](#) into the clotted artery and inflating it to dilate the artery. A fine mesh of wire, called a stent can additionally be placed inside the artery as a permanent endoprosthesis to reduce the risk of future [restenosis](#). The main steps of these procedures are illustrated in Fig. [1.6](#). A more detailed description of this therapy will be given in section [1.2](#) since it is the procedure of interest for our work.

1.1.6 Coronary artery bypass graft (CABG)

CABG is a surgical procedure performed to relieve angina and reduce the risk of death from coronary artery disease. Arteries or veins from elsewhere in the patient's body are grafted to the coronary arteries to bypass atherosclerotic narrowing and improve the blood supply to the coronary circulation (Fig. [1.7](#)).

1.1.7 Risk factor Management

None of the exposed treatment options actually removes or reduces the plaques that lie inside the arteries. In the case of the CABG they are bypassed and in the case of the angioplasty they are pushed away from the lumen. If the risk factors are not taken into account by the patient it is highly likely that the arteries will get clotted again. The important risk factors fall into two categories, the uncontrollable ones, and the controllable ones. In the first category fall age, (male) sex and heredity. The second




FIGURE 1.6: Balloon angioplasty [NIH10b] and stenting [NIH10c] procedures.




FIGURE 1.7: Illustration of a bypass [NIH10a].

category includes disease and disorders that shall be treated such as hypertension or diabetes and many lifestyle related factors: smoking, low fat/salt diet, exercising and avoiding chronic high stress levels for instance.

1.2 A more detailed view on stenting

So far, the general clinical context has been exposed. We propose now to get into more details on interventional cardiology and stenting. Some key figures about this field are presented in table 1.1.

Diameter of healthy coronary arteries	$2 \leq .. \leq 4, 5$ mm for the sections of interest
Coronary stent diameter [GS10b]	$1, 5 \leq .. \leq 4, 5$ mm
Coronary stent length [GS10b]	$10 \leq .. \leq 40$ mm
Diameter of a stent strut [GS10b]	0.1 mm
Contrast to noise ratio of a stent [RLM+05]	$0.19 \leq .. \leq 1.3$
Xray detector pixel dimension	0.2 mm
Conversion factor	0.14 mm/pixel
Cost of an angioplasty [FYL+03]	\$12, 700
Cost of a drug eluting stent [MFKY11]	\$3000
Cost of an IVUS catheter [Wik11b]	\$600
Cost an FFR pressure wire [FYL+03]	\$550
Number of stents sold in the world	some millions / year
Speed of a coronary artery	$2 \leq .. \leq 11$ cm/s
Heart beat frequency	1 beat/s
Duration of an angioplasty (except CTO)	$20 \text{ min} \leq .. \leq 1 \text{ h}$
Xray dose of a coronarography [Mac09]	57 Gy.cm^2 , 1270 images
Xray dose of an angioplasty [Mac09]	94 Gy.cm^2 , 1355 images
Complication rate [SJP+09]	3%

TABLE 1.1: Some key figures about PCI (typical orders of magnitude).

1.2.1 A short presentation of interventional cardiology

A very short historical perspective [GS10b] on interventional cardiology will enable us to introduce all the concepts necessary to understand the technique as it is performed nowadays (more information about this field can be found in [RCT10]). Interventional cardiology is a young, successful, rapidly changing field. The history of interventional cardiology has been forged by repeated loops consisting in identifying the major limitations of the current technique and developing technological innovations to overcome them. The main steps of this recent history, that build the clinical context of the current PhD thesis are the following. First Angiography in 1964 by Charles Theodore Dotter and Melvin P. Judkins: contrast media was injected into vessels to make them visible on radiographs (Fig. 1.8 and 1.9). In 1977 Andreas Gruntzig performed the first balloon coronary angioplasty, it was the birth of interventional cardiology (Fig. 1.6). However Plain Old Balloon Angioplasty (POBA) was hindered by the problems of acute vessel closure and restenosis. Coronary stents developed in the mid 1980s replaced POBA. In 1986 Sigwart et al. implanted the first coronary stent (Fig. 1.10). It reduced the rate of emergency coronary artery bypass. However subacute thrombotic coronary occlusion was observed in 18% of the cases within two weeks of implantation. Progress included dual antiplatelet therapy and care for good deployment that made the technique safe (studies BENESTENT [SKD+99] and STRESS [GBB+98]). In addition to the risk of subacute thrombosis, another problem araised as a consequence of stenting called *in stent neo intimal hyperplasia*. It is the intrastent growth of scar tissue resulting in restenosis rates of 20% to 30% (Fig. 1.11). In order to solve this, Drug Eluting Stents (DES) have been developed with great success in the early 2000s. Clinical trials (summarized in

[GS10b]) demonstrated the superiority of DES over [Bare Metal Stents \(BMS\)](#) resulting in significant reductions in angiographic in-stent-late-loss, in-stent angiographic binary restenosis, and repeated vascularisation at both short term and long term follow up, over numerous types of patients and lesions. However rates of death and [myocardial infarction](#) at long term follow up are comparable. Today 75% of the revascularisation procedures for patients suffering from CAD are performed with DES.




FIGURE 1.8: [Angiograms](#) of Dotter's 1st [catheter](#) patient, Laura Shaw: A) before transluminal dilation of the left superficial femoral artery, B) immediately after dilation, and C) 3 weeks after the procedure (image from [Pay01]).




FIGURE 1.9: Visualization of a right coronary artery under [Xray fluoroscopy](#) before and during contrast media injection. We note that in the absence of contrast the medical tools inside the artery are visible but not the artery itself. When contrast media is injected, the artery is visible but the medical tools become hardly visible (images obtained on an Innova 2100 system of GE healthcare).

1.2.2 The cathlab

The PCI procedures take place in a dedicated interventional room called [catheterization lab](#), or simply cathlab. At the center of the cathlab is the imaging system, in our case a system from the Innova series manufactured by GE (Innova 2100, or 2100^{IQ} for instance [Ele11]). The imaging of the coronary arteries (coronarography) and the deployment of stents (PCI) are minimally invasive techniques monitored under [Xray fluoroscopy](#) with




FIGURE 1.10: Various popular coronary stents. From left to right: Cypher [Hos07], Endeavor [ptc09], Taxus Liberte [DC09], Xience V [DC10b]. Note that the metal meshes are quite different from one stent to another and can be rather complex.

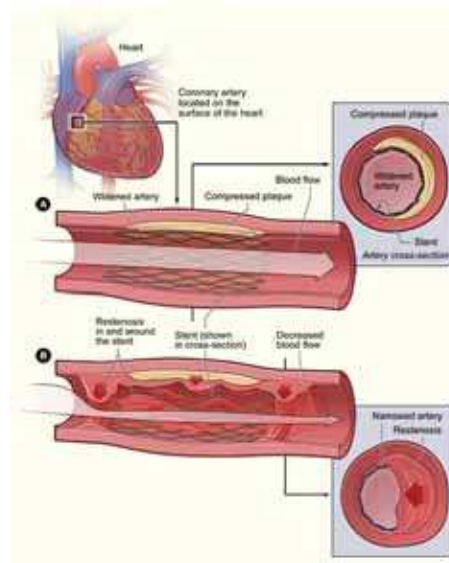


FIGURE 1.11: In stent restenosis (image from [NIH10d]). This complication was common with bare metal stents. It decreased very significantly with the introduction of drug eluting stents.

the Innova system. At the demand of the clinician, real-time Xray video is produced at 7.5, 15 or 30 images per second. Any part of the body can be imaged, under virtually any angle, thanks to the positioning capabilities of the system: a moving table (3 degrees of freedom in translation), on which the patient lies, and an imaging chain (3 degrees of freedom in rotation) that can be positioned around the patient. The images are displayed on suspended monitors, right in front of the clinician. Fig. 1.12 shows the Innova system, and the positioning of the patient, the clinician and the screens. The system also offers the possibility to spin the image chain around the patient, in order to perform 3D reconstructions for instance. There are two main types of image acquisition commonly available: cine (or graphy) and fluoro (also called fluoroscopy, or scopy). The cine images are acquired at a dose level 5 to 10 times superior to fluoro. Therefore they have better image quality. They are automatically stored on the system and are used for diagnosis and documentation of the stages of the treatment. Fluoro images, of lower image quality and dose, are used for the longest parts of the intervention to monitor the medical tools inside the patient body. They are not systematically stored. In some cases, when cine

is required to document a given stage of the intervention, it may be replaced by fluoro when it is stored (on request of the clinician). The appropriate usage of both and of the multiple dose customization capabilities of the system enables performing interventions according to the ALARA principle that rules the image quality/dose trade-off: Xray dose must be As Low As Reasonably Achievable.

The technical developments presented in the following chapters are typically designed to be potential features of cathlab systems.



FIGURE 1.12: Innova system - GE Healthcare

1.2.3 Breakdown of a stenting procedure

The stenting procedure follows a set of pre-defined steps. We detail them, and for each of them express what are the challenges met by the clinician. We refer to this breakdown in several chapters to explain which clinical need we are addressing, and how the workflow is possibly modified by the proposed solution.

In practice, the interesting part of the procedure starts when the clinician inserts an injection catheter at the entry of the coronary artery to produce reference images of the pathology. These images, illustrating the vessels, help the clinician decide upon the treatment strategy. If stenting is selected, he/she will keep this reference image on one of his screens to help him guide his gesture all along the procedure. The imaging chain typically stays in the same position until the end of the treatment phase. Then a [guide-wire](#) is inserted down to the distal end of the artery. The crossing of the lesion may present some difficulties depending on its characteristics. This guide-wire will serve as support to slide all the necessary medical tools. A balloon is slid down to the lesion and inflated at high pressure (typically 12atm). Then a stent, mounted on a balloon is slid down to the lesion and inflated. Finally, the balloon is removed and a record acquisition is performed with contrast injection to verify that the lesion is properly treated. To complete the check the same injection is performed again in a perpendicular

	Step	Acquisition	Clinical challenge/unmet need
1	Reference injection	C	Assessing if the lesion shall be treated; choose balloon and stent dimensions
2	Inserting guide-wire down to the end of the artery	F	some lesion/bifurcations can be difficult to cross
3	Sliding balloon down to the lesion. Contrast injections.	F	Visualize the position of the balloon versus the environment
4	Balloon inflation	F or C	
5	Removing balloon	F	
6	Contrast injection to assess the lesion	F or C	
7	Sliding stent down to the lesion. Contrast injections.	F	
8	Balloon inflation / stent expansion.	F or C	Assessing stent deployment
9	Balloon removal.	F	
10	Contrast injection	C	Assessing if the lesion is treated and if the stent is correctly deployed
11	If the lesion is not properly treated repeat step 8	F	
12	Control injection in perpendicular angulation	C	Assessing if the lesion is treated and if the stent is correctly deployed
13	If the lesion is not properly treated repeat step 8	F	

TABLE 1.2: Steps of a stenting procedure. First column is the index of the step, second one its description, third one the acquisition mode (F stands for fluoro and C for cine), and the last one the clinical challenges. Steps 3 to 6 can be omitted in case of direct stenting.

angulation to account for the 3D nature of the lesion. If necessary a balloon is brought back and further inflated again. The main steps of a stenting procedure are reported in Table 1.2.

1.2.4 On the importance of assessing stent deployment and apposition

Many clinical studies demonstrated that improper stent sizing, deployment or apposition have a strong impact on the success of the procedure. The main possible resulting complications are restenosis, the need for target vessel re-vascularisation, stent fracture and Stent Thrombosis (ST). The figures reported in this section that support these assertions are all extracted from [GS10b].

- **Stent Thrombosis:** One of the major safety concerns with stenting in today's clinical practice is ST and especially very late ST (more than one year after stent implantation). The annual risk rate of very late ST is between 0,36% and 0,6% and persists for at least 5 years after implantation with DES. Although very rare, this event is critical because it has very poor prognosis: 10% to 30% of patients with definite ST will die. Among the possible precipitants of stent thrombosis we retain a subset that we will be able to interact with in the work presented here:

calcification of the lesion, stent undersizing, inadequate stent expansion, residual stent stenosis and incomplete stent apposition. Although ST re-entered under the spotlight in the DES era, it can also happen with BMS and is also associated to unoptimal PCI results.

- **Stent under-expansion or inadequate expansion:** We have just seen that it is related with complications. It is a phenomenon of particular interest for our work since it is typically the kind of finding that we hope the techniques we will present will enable. It has been shown to be associated with increased rates of restenosis, target vessel re-vascularisation, stent fractures, and to be a possible precipitant of ST. The rate of stent under-expansion with DES may be as high as 30%. If it can be detected during the stent deployment procedure it is recommended to maximize the final stent area/diameter with a (non-compliant) balloon inflation as a corrective action.
- **Stent fracture:** It is an uncommon complication after DES whose true incidence remains unknown, ranging from 1% to 30% depending on the studies. However it has been shown that 70% to 80% of patients with a stent fracture will present in-stent restenosis or ST. If detected, some action can be performed, mainly repeated PCI.

These facts emphasize on the utility of visualizing the stent and assessing its deployment in real time in the cathlab at the time of the procedure. This is the core clinical need addressed by all the image processing applications that are presented in the following chapters of this thesis.

1.2.5 A word about calcifications

Some lesions (stenoses) when aging can develop calcific layers (Fig. 1.13). This aspect of the lesions is checked by the clinicians and is taken into account in the American College of Cardiology/American Heart Association (ACC/AHA) stenosis morphology classification (making the lesion at least of type B). This calcium happens to complicate the stent deployment since it makes the vessel more rigid. It can impact the choice of the tools (the balloon and its pressure) used for (re)deployment and of the procedure since the clinician may pre-dilate before stenting to maximize the chances of proper lesion treatment. Visualizing the calcium and its position relatively to the stent is thus of interest to the clinician. Moreover calcium deposit has gained interest over the past years in the context of CT scans since it has been demonstrated to provide predictive information beyond that of classical risk factors [GLA⁺04].



FIGURE 1.13: Calcification are very well visualized with Computed Tomography [DC10a]. They appear in white, around the lumen of the artery.

1.2.6 Synthesis of the needs of the clinician in PCI

According to the clinical context we have exposed, we are now in position to synthesize the needs of clinicians in PCI. These needs are different according to the stage of the procedure. Before expanding the stent (step 6 of table 1.2), the clinician is interested in the exact positioning of the balloon versus the lesion, the calcifications, the potential bifurcations and potentially pre-existing stents from prior procedures. Right after stenting (step 8 of table 1.2), the clinician is interested in the shape of the stent to assess its proper expansion/deployment and its relative position to the vessel wall. Quantitative assessment of stent expansion (measurement of diameters) is also relevant in this context. At patient follow-up, detecting potential stent fractures (that typically appear late after stent deployment as a result of fatigue) is one of the concerns. Therefore a cathlab shall ideally provide:

- visualization of the position of the balloon with respect to the vessels, stents and calcifications,
- visualization of deployed stents,
- visualization of the relation between deployed stents, vessels and calcifications, and
- quantitative measurements of stent deployment.

1.3 Stent visualization in the cathlab

Let us now review how conventional imaging techniques address the needs we have expressed with respect to PCI. There are two families of imaging techniques in the cathlab: Xray imaging with the interventional system as presented so far, and intra-vascular imaging, with ultra-sound or optical coherence tomography. The properties of the stents

are such that they have very low visibility in conventional Xray images. Therefore, intravascular imaging is the preferred option to visualize them. To better understand this, let us consider more precisely what a stent is. A look at the data sheets of the most common stents (Cypher, Taxus Express, Taxus Liberte, Endeavor and Xience V) shows that they share the following characteristics [GS10b]: they are made either of stainless steel or CoCr, the strut thickness is between 81 and 140 μm , and they are fine complex meshes. The SNR under Xray is directly related to the material and its thickness. Although metals have great radio-opacity, the struts are so fine that the resulting attenuation of the Xray beam is minor. Simulations [RLM⁺05] show that for typical patient thicknesses the contrast to noise ratio of the stent is between 1.3 and 0.19. In addition to these discouraging figures, the very nature of the geometry of the stent (Fig. 1.10), a complex 3D mesh, adds to the difficulty of assessing it from projection images. Conversely in intra-vascular imaging the stent is visualized into cross sections. Therefore its geometry is simple to understand, and each strut can be visualized separately. Moreover its contrast to noise ratio is such that strut conspicuity is high, and that they can be immediately identified in the images. Finally the vessel is also visible in the cross sections, enabling the clinician to assess the apposition of the stent onto the vessel (Fig. 1.14). However, it is important to note that intra-vascular imaging is a significant additional cost over the regular procedure and is therefore only moderately used (at the exception of Japan where it is reimbursed according to the national policy). Moreover, there is a large number of vessel slices to examine, and it is not mentally straightforward to build an image of the general stent scaffolding out of them.

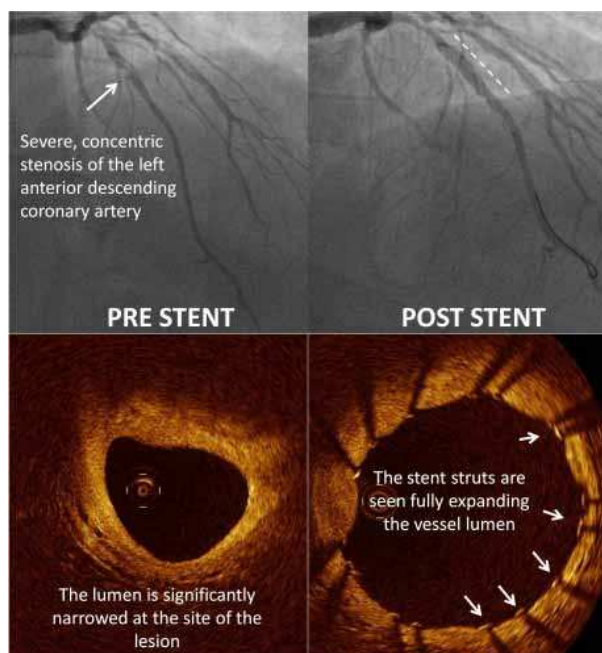


FIGURE 1.14: Stent cross section in OCT versus conventional angiography [Lod10].

In short, conventional Xray imaging does not allow to visualize properly the stent and intravascular imaging, although better suited, is longer, more complex and costly. There is no imaging techniques providing a cheap, convenient and easy way to visualize stents inside the cathlab during every procedure.




FIGURE 1.15: Coronary stent image with conventional Xray (image zoomed on the stent). Note that the stent is almost invisible, only the guide-wire and the two balloon markerballs are clearly visible (image obtained on an Innova 2100 system from GE).

1.4 Technical context

1.4.1 Interventional device detection and tracking

Over the past years growing interest has emerged for the real time or near real time detection and tracking of interventional tools in Xray fluoroscopy. Several teams aimed at detecting the balloon marker balls [SFL⁺09a, FNLR08, RLM⁺05, KL05], the guide-wires [BAG⁺07, HVP10b, HVP10a, HGG⁺09, BVG09] and catheters [BLHS09, BWL⁺10] with different applications in mind : 3D/2D registration [BVG09, GBK⁺09], motion compensation [BVG09, BLHS09, GBK⁺09], navigation [GBK⁺09, BWL⁺10], stent visualization [SFL⁺09a, FNLR08, RLM⁺05, KL05], guide-wire enhancement [BAG⁺07, HVP10b, HVP10a, HGG⁺09], denoising [SSD⁺05, BV08a], 3D reconstruction of stents, arterial segments and closure devices [SFL⁺09a, Per08]. This non exhaustive list illustrates this general trend. The image processing algorithms developed in this thesis fall into this context. They are all related to the near real time detection, tracking and registration of markerballs and guide-wires in order to enhance the visibility of stents.

1.4.2 Digital stent enhancement

These algorithms mostly fall into a category that we will call Digital Stent Enhancement (DSE). The cornerstone of DSE is an idea that emerged in the early 2000 [CAW00,

[CAW02, CWA03, FNLR08]. It is based on the observation that the main limitation on the SNR of stents is the limited Xray dose delivered in each frame. However, along an image sequence the whole dose that is delivered would be sufficient to image the stent with a much higher contrast to noise ratio (typically 5 times better than in a single frame). Consequently DSE proposes to track the stent along the image sequence and to integrate the value of the pixels of the stent along their trajectories. For each stent pixel, all the information present in each image of the sequence is cumulated. As a consequence, in the resulting image the stent is greatly enhanced. Since the stent itself is extremely hard to detect, its motion is inferred from the position of the balloon markerballs and of the guide-wire.

The full automation of the DSE algorithms is key to their adoption by clinicians in daily practice. Indeed, during the course of an intervention the clinician is extremely busy: his hands manipulate the interventional tools, his feet control the X-ray on/off pedals and his eyes focus on the display of the imaging system. Any user interaction that distracts his attentions from these tasks is considered a serious limitation.

1.4.3 Context of the thesis

I performed the work presented in this PhD thesis concurrently with my full-time research-engineer position at GE Healthcare. I had held this position since nearly 6 years when I started this project, in collaboration with the A3SI team at ESIEE. Both institutions were complementary: GE impuled a result-oriented and application-driven direction to my work, while ESIEE was instrumental in the formalization of my contributions and in the developement of new techniques. From an organizational standpoint, I dedicated one day per week to my PhD activities at ESIEE, and the rest of the week to my research-engineer position. The time spent at GE was mainly dedicated to the developments of image processing algorithms, their evaluation in clinical setting and their adaptation into industrial products. The time at ESIEE was spent on brainstorming, writing, learning and teaching.

The philosophy of this PhD thesis is to develop image processing applications to support clinicians in their daily practice. Some of the developments presented in this manuscript (Chapters 2 and 3) have recently been implemented into commercial products (*StentViz* - GE) and are used in daily practice by clinicians (Fig. 1.16). To this end, I performed the following tasks¹ listed in chronological order:

- develop the algorithm,
- install a prototype on a clinical site and improve it based on clinical feedback²,

¹under the supervision of Régis Vaillant - Principal engineer.

²This work has been done in the cardiology department of Institut René Dubos (Pontoise - France), under the supervision of Dr. F. Funck.

- lead the clinical validation study on the GE side,
- issue patents protecting our intellectual property and publications to valor our technology,
- deliver the implementation of the industrial product,
- define the specifications and tests of the application,
- deliver design, specification and test documents,
- demonstrate substantial equivalence with an existing industrial product for the approval on the US market by the FDA.

I benefited during this work of a hardware/software platform developed internally that allows to seamlessly install prototypes on sites without impacting the clinical workflow. In order to reach the expected level of performance for an industrial product, we paid particular attention to the validation of the performance of each algorithmic brick on large sets of clinical data (from 20 to 200 clinical sequences depending on the chapter). Up to now, we keep on monitoring the feedbacks associated with these products from the hundreds of equipped sites.

The technical developments presented here address a large number of topics starting with feature detection and tracking (Chapters 2 and 4), passing through registration, curvilinear structure segmentation (Chapter 5) and layer separation (Chapter 3). I have chosen to present this PhD guided by the applicative direction that is common to the entire document: device visualization enhancement. It has been the subject of several publications and patent applications [BV08b, BVG09, BVF⁺11, VLB06, BV10, VB10, FGBV09, FLBV10].

1.5 Structure of this thesis

The chapters of the thesis address DSE and related topics in the following manner:

- **Chapter 2: Stent visualization enhancement algorithm.** This key chapter presents a thorough analysis of the problem of DSE, a fully automated and image quality oriented algorithm to solve it, detailed clinical validation and comparison with an existing commercial DSE software.
- **Chapter 3: Subtraction of the guide-wire in enhanced stent images.** We take advantage of the transparent nature of Xray images to overcome a common limitation of DSE: the guide-wire, that has no clinical interest, overlaps with the stent and hides it. The technique we propose, allows to erase the guide-wire from the images without degrading the subtle details of the stent struts.



FIGURE 1.16: Clinicians performing an angioplasty. Note the use of the stent visualization enhancement software on the top left screen that is the output of the work presented in Chapter 2. This image depicts a stent that has just been deployed with an image quality far superior to conventional Xray [FGBV09]. Image courtesy of Dr. T. Lefèvre, Dr. M-C. Morice, Dr. T. Hovasse, Dr. B. Chevalier, Dr. Y. Louvard - Institut Cardiovasculaire Paris Sud, Massy, France

- **Chapter 4: 3D stent reconstruction.** Based on ideas similar to DSE, a 3D stent reconstruction technique has recently been proposed, that requires detecting and tracking the balloon in a spin acquisition. Given the length of the spin sequences, and the changes of appearance of the balloon under different viewpoints, the tracking developed in the previous chapters had to be largely revisited. We present a fully automated balloon tracking algorithm dedicated to 3D stent reconstruction and assess its performance on clinical datasets. Moreover we extend the concepts of non-linear registration and guide-wire subtraction of the previous chapters to the case of 3D stent reconstruction.
- **Chapter 5: Curvilinear structure enhancement.** We investigate here a new approach to curvilinear structure enhancement based on the combination of local shortest paths and voting. In a first part, curvilinear structure enhancement techniques and evaluation methodologies are presented. Then our technique is described and the links with other existing techniques are discussed. Finally we evaluate its performance for the enhancement of guide-wires in Xray images.

Each chapter is organized as a independent publication (abstract - introduction - background - method - result - discussion - conclusion) and can be read quite independently from the other ones. However there is a logic in the organization of this thesis and the reader will benefit from reading them in order: the second chapter introduces the problem of stent visualization and our digital stent enhancement technique. In Chapters 3 and 4, we elaborate on it. Chapter 5 presents a more standalone work dedicated to the enhancement of curvilinear structure in images. It is applied to the detection of the guide-wire, a problematic discussed in the 3 previous chapters.

1.6 Main technical contributions

The most significant contributions of this thesis can be classified into the following topics: stent visualization, interventional tool detection and tracking, curvilinear structure enhancement and methodology.

1.6.1 Stent visualization

A large proportion of the contributions of this thesis deal with stent visualization in the cathlab. We propose several technical elements that are new to this field. We quantify their performance and demonstrate them on clinical cases.

State of the art registration techniques in this domain, that handle translation, rotation and scaling of the stent, are linear. In Chapter 2, we introduce a non-linear one especially tailored to the specificities of stents that takes into account their *a priori* mechanical properties. It can handle deformations of the stent associated to complex vessel motion, like bending for instance. We demonstrate the superiority of this model (Fig. 1.17) versus the classical one on a large database of cases in a quantitative manner. This work has been presented at conferences [BV08b, FGBV09], published in a journal [BVF⁺11] and patented [VLB06].




FIGURE 1.17: Stent visualization enhancement results obtained with: (a) state of the art linear registration, (b) our non-linear one. Observe that the stent is sharper in (b).

To further improve the visualization of stents we develop in Chapter 3 an original approach that erases the guide-wire from Xray images. It takes advantage of the transparency of the images acquired with this modality. We express the guide-wire removal task as a layer separation problem where the image is modeled as a sum of three layers: the background, the guide-wire and the stent (Fig. 1.18). The first two layers are estimated and subtracted from the original image to isolate the stent layer (Fig. 1.19). The performance of this technique has been demonstrated on a set of clinical cases. Results

have been presented at a conference [FLBV10] and some technical aspects have been patented [BV10, VB10].




FIGURE 1.18: Planar layer modeling of a DSE image: 3 separate layers, namely the anatomical background, the guide-wire and the stent, are combined in the image formation mechanism.




FIGURE 1.19: Left: input image. Center: DSE image (note that the background layer has been removed). Right: guide-wire subtracted image.

Our contributions in stent visualization also encompass 3D reconstruction of stents (Fig. 1.20). In Chapter 4, we extend the tools originally developed for the 2D case: non-linear registration and guide-wire subtraction. We demonstrate how they can be naturally adapted to this framework by processing the 2D projection images prior to the reconstruction (Fig. 1.21 and Fig. 1.22).

1.6.2 Interventional tool detection and tracking

The 2D and 3D stent visualization techniques that we propose, rely on the automatic detection and tracking of two interventional tools: the radio-opaque markers embedded in the angioplasty balloon and the guide-wire supporting it. In Chapter 2, we describe an algorithm to perform this task in static acquisitions (Fig. 1.23) and we adapt it to the rotational case in Chapter 4 (Fig. 1.24) by revisiting the cinematic hypotheses and taking into account 3D motion consistency. These fully automated methods yield state of the art performance.




FIGURE 1.20: 3D coronary stent reconstruction at a bifurcation.




FIGURE 1.21: A 3D stent slice: on the left classical processing, on the right the guide-wire has been subtracted between the markers. The stent border that is hidden by the guide-wire on the left is visible on the right, at the location indicated by the arrows.




FIGURE 1.22: Volume rendering of a reconstructed stent. Left: state of the art reconstruction technique. Right: The reconstruction benefited from the extension of our stent techniques to the 3D case. The motion compensation is performed with non-linear registration and the guide-wire has been removed from the image. Notice that the overall stent shape is easier to delineate.



FIGURE 1.23: Left: a frame of a short static acquisition. Center: automatically segmented markers and guide-wire. Right: Close up on the segmented interventional tools.




FIGURE 1.24: Markers and guide-wire automatically segmented in three frames of a rotational acquisition (top row: original images - bottom row: segmented interventional tools). Observe that the markers exchange their relative positions as a result of the gantry rotation.

Moreover, we propose in Chapter 2, a technique to segment a guide-wire between two given points. The main challenge in this application is to design a technique that is particularly robust to noise. State of the art techniques usually rely on the computation of a geodesic path between the two end points. The geodesic minimizes the integral of a potential along the path. Our approach differs by two aspects:

- First, we define a metric based on the shape of the integrated profile along the path (Fig. 1.25). If the path is aligned with the guide-wire, the integrated profile shows a characteristic shape with a clear minimum at its center. Otherwise the

minima of the profiles sampled along the path do not add up coherently and the integrated profile does not present the expected shape.

- Moreover, contrarily to classical shortest paths approaches, the paths are constrained to be smooth parametric curves (piecewise polynomial). This prevents them from wandering around the actual guide-wire under the attraction of local noise fluctuations (Fig. 1.25).

We demonstrate the ability of our technique to handle large amounts of noise on clinical data (Fig. 1.25). It has been patented [VLB06] and published in a journal [BVF⁺11].




FIGURE 1.25: Top left: an extract from a cine image depicting a set of smooth paths between two endpoints. Top center: A clinical image corrupted with a high level of noise. Top right: guide-wire segmented with our method between the two markerballs. Bottom left: A clinical image depicting a guide-wire. Bottom center: 3 candidate curves (note that the blue one is well aligned with the guide-wire centerline whereas the red and green ones are not). Bottom right: The three averaged profiles (same color code as bottom center). Abscissa is the signed distance to the curve and ordinate is the pixel intensity. Only the profile of the blue curve that is aligned with the guide-wire exhibits a marked minimum at its center.

1.6.3 Curvilinear structure enhancement

We present in Chapter 5 a new tool to process curvilinear structures in images: the polygonal path image (\mathfrak{P}). It is based on the concept of locally optimal paths and has several interesting properties:

- a unification of local, semi local, and global curvilinear structure analysis in a single framework,
- the ability to control the smoothness and length of the structures to analyze,
- an intuitive and simple parametrization,
- an efficient computational scheme.

\mathfrak{P} is a rich descriptor of the curvilinear structures present in an image from which we derive line enhancement and direction field computation techniques. We have demonstrated the relevance of this line enhancement technique for the application of guide-wire segmentation both qualitatively (Fig. 1.26) and quantitatively using the ROC analysis formalism. We show that it outperforms previous state of the art techniques. Moreover, contrarily to the direction fields built with state of the art methods that are not relevant outside of the curvilinear structures, \mathfrak{P} indicates the direction of the closest such structure (Fig. 1.26).

1.6.4 Methodological contributions

We finally want to emphasize on the significant and systematic efforts that we dedicated to the evaluation and validation of the proposed algorithms in clinical settings. The image quality of the digital stent enhancement technique studied in Chapter 2 has been validated on nearly 200 clinical cases and its detection and tracking capabilities on several thousands of frames. Similarly, the detection and tracking algorithm in rotational acquisitions (Chapter 4) has been validated on a set of 22 sequences totalizing 3000+ frames. The performance of our line enhancement technique has been quantified on 100+ guide-wire images. Moreover, each validation covers the design of assessment metrics. One of the challenge of this task is to translate the performance expected by the user into measurable technical quantities. We believe that these methodologies and their associated results can be considered as best practices and constitute a relevant comparison point for other researchers in our field.

1.6.5 Material published for the first time

Several of the studies presented in this manuscript appear publicly for the first time. Among these, we would like to point out:

- The comparison of our stent enhancement technique to an industrial product (Section 2.4.5)
- Most of the content of Chapters 3 (guide-wire subtraction) and 4 (3D stent visualization).




FIGURE 1.26: Illustration of two applications of \mathfrak{P} . Top left: a clinical image illustrating a guide-wire. Top right: line enhancement by smooth local path voting. Bottom left: aggressive contrast setting on the image displayed in top right reveals its structure: A path attracted by linear structures originates from each pixel of the image. Bottom right: Direction field estimated with \mathfrak{P} .

- The entire content of Chapter 5 (locally optimal path technique for curvilinear structure processing).

1.7 The future of the cathlab in interventional cardiology

Before concluding this introductory chapter, let us have a look at what we can foresee of the near future of the cathlab for PCI related procedures. In the 90s and early 2000 the introduction of BMS and DES has been a complete game changer in cardiology enabling to treat safely a very wide population of patients. Nowadays, the trend is to improve the short and long term safety of re-vascularisation procedures. It includes, on the one hand, doing it better (safer stents, additional tools) and performing it only when it is best suited (improvements in decision making).




FIGURE 1.27: Top [Med11] principle of the FFR: the device allows to compute the ratio of pressures distally and proximally to the lesion in the hyperhemic state. Bottom [CJHJ05] effect of thrombus aspiration and balloon dilatation in an occluded artery.

Stent design is an extremely active research area [GS10a]. Current developments focused on stent safety include reducing the thickness of the struts and developing biodegradable stents [GS10a]. Both of them raise issues concerning the radio-opacity of the new stents, and will probably raise the level of expectations on DSE algorithms.

Besides, two recent studies may impact the “stenting procedure breakdown” that we have presented (see Fig. 1.27). The first one [SVvdH⁺08] demonstrates that adding a step to this procedure is beneficial in case of acute MI. Instead of proceeding to stenting and potentially crushing the thrombus onto the vessel wall, the emerging practice tends to perform thrombus aspiration with a specific device prior to stenting. The second study [TDBP⁺09] deals with the choice of the lesions to stent. They propose to systematically base this choice on a measure of the hemodynamic impact of the lesion, the Fractional Flow Reserve. This measure is also performed with a specific catheter mounted device. It will probably result in less lesions being stented.

Conversely, PCI has conquered a new territory over the past years. **Chronic Total Occlusions (CTO)** that used not to be treated, are now more and more handled by PCI. These complex and often long procedures will probably drive new image processing applications based on clinical tool detection and tracking. From a more prospective standpoint, current research efforts aim at curing the consequences of MI and restoring valid heart tissue. These techniques based on **stem cells** injection and **angiogenesis** are still at a very early stage.

1.8 Conclusion

This thesis presents stent imaging techniques based on conventional Xray acquisition that aim at addressing previously unmet needs of the clinicians. To this end, we present a set of image processing algorithms, that typically incorporate detection, tracking and registration of the clinician's tools. The resulting techniques provide relevant clinical information that was not accessible in the original X-ray images. Moreover it does not increase the time or cost of the procedures. From a long term perspective, the development of algorithms dedicated to the processing of the clinical tools in interventional imaging is a trend that is likely to increase and find new applications with emerging procedures, ultimately improving device guidance with real-time implementations.

Chapter 2

A stent visualization enhancement algorithm

2.1 Abstract

In this chapter¹ we propose a comprehensive study of Digital Stent Enhancement (DSE), from the analysis of requirements to the validation of the proposed solution. First, we derive the stent visualization requirements in the context of the clinical application and work-flow. Then, we propose a DSE algorithm combining automatic detection, tracking, registration and contrast enhancement. The most original parts of our solution: landmark segmentation and non-linear image registration are detailed. Finally, we validate the algorithm on a large number of synthetic and clinical cases. Performance is characterized in terms of image quality, automation, and execution time. This work is, to the best of our knowledge, the first comprehensive work on DSE, covering problem statement, proposed solution, and validation strategies.




FIGURE 2.1: Left: a coronary stent imaged with Xray. Right: Result after DSE.

¹The technical work we present in this chapter has been partly published in a conference [BV08b] and a journal paper [BVF⁺11]. The key points of the technique are covered by a patent application [VLB06] and some of the results were the subject of an oral presentation at EuroPCR 2009 and have been published in an abstract [FGBV09]. Finally, the comparative data of Section 2.4.5 have supported the pre-market notification of StentViz (GE) [Mor09].

2.2 Introduction

The general clinical context of this thesis has been extensively described in Chapter 1. In order to keep this chapter self-complete, we start with a brief reminder of the most important elements necessary to understand the covered topic. Coronary heart disease is the most common cause of sudden death, and the most common reason for death of men and women over the age of 20 [RFF⁺07]. This disease tends to narrow the lumen of coronary arteries by the accumulation of atheromatous plaques within their walls. The narrowing of the artery lumen due to plaque progression, or plaque rupture can be cured by the expansion of a fine metallic mesh called a stent that is implanted in an artery wall acting as a scaffolding to open the lumen thereby restoring blood flow. Stent placement is performed as part of a percutaneous coronary interventional procedure. The procedure is typically performed under the guidance of X-ray fluoroscopy delivering real time video of the clinical tools and devices in the patient's anatomy, (Fig. 2.2 (a)), typically at 15 fps. The clinician first introduces a metallic guide-wire inside the artery that serves as support for sliding an angioplasty balloon equipped with a stent. The angioplasty balloon is used to expand the stent. Its inflation simultaneously opens the lumen and embeds the stent into the vessel wall. In order to visually assess the location of the balloon/stent on the guide-wire, the balloon is equipped with two highly radio-opaque marker-balls delimiting the position and extent of the devices. Fig. 2.2 (b) depicts a part of an X-ray image, with a guide-wire, the deflated balloon (invisible), its two marker-balls and the deployed stent. Proper positioning of the stent and apposition onto the vessel wall are key to the success of the procedure and patient safety [GS10b, ASPV⁺07, FCM⁺05, BDJ⁺04].

However visualizing stents with conventional X-ray images is extremely challenging. Indeed, coronary stents are characterized by their low radio-opacity and their fast motion. The recommended solution to assess stent deployment and apposition is to perform Intra-Vascular Ultra-Sound imaging (IVUS). This modality, considered as the gold standard for stent and vessel assessment during percutaneous coronary interventions, provides cross sectional images of the vessel, depicting the stent, the vessel wall and enables quantitative length and area measurements. Unfortunately systematic use of IVUS is impractical since it would add significant time and cost to the procedure. It is only used in 4.5% of the cases in Europe, 14% in the USA and 60% in Japan where it is reimbursed. In this chapter, we consider techniques enabling stent visualization that are consistent with the normal workflow of an angioplasty procedure. Although they cannot bring direct information on the apposition of the stent onto the vessel wall, they can provide relevant images for assessing deployment irregularities, lesion treatment and potentially measuring stent expansion.




FIGURE 2.2: (a) An image from a clinical image sequence with the clinician’s tools positioned in the patient’s vasculature; (b) the region in (a) delimited by the gray square with annotation of tools/devices; (c) the region in (b) inside the gray square with a deployed stent; (d) demonstration of the improved stent visualization by the application of DSE to (a). This last image has been produced with StentViz, the DSE feature of GE, that is partly described in this chapter.

Since 2000, an image processing technique called digital stent enhancement (DSE) has emerged to produce an enhanced image of a stent from an X-ray image sequence [CAW00, FNLR08, BV08b]. It has gained interest from clinicians over the past years, leading imaging system manufacturers to collaborate with them to study this topic [Koo05, RLM⁺05, KL05, CSMS05, MVP⁺07, OSB⁺08, BV08b, CAC⁺09, FGBV09, ZHB⁺08]. To perform DSE, the clinician first shoots an X-ray sequence of the deployed stent while keeping the delivery balloon in place. It is typically a short X-ray sequence showing the stent at 15 fps where the stent motion is induced by the heartbeat and the breathing of the patient. As in any regular X-ray sequence, the visibility of the stent is limited. The motion of the stent is inferred by the DSE algorithm based on the motion of balloon marker-balls that serve as landmarks. Motion compensated image integration is the corner stone of DSE and enables a significant reduction of the noise while preserving the details of the motion compensated components of the image (Fig. 2.2 (c) and (d)). Studies show that stent visibility is significantly improved by this technique [Koo05, OSB⁺08, FGBV09]. Moreover this image quality improvement fosters the ability of an observer to detect stent under deployment [RLM⁺05].

DSE is thus a special case of motion compensated noise reduction, a temporal filtering technique commonly applied to enhance multimedia image sequences [Mar87, OST93, Sam86, SOF91, DS84, BKE⁺95, SL03, ZPP04, BK95]. The application of such techniques to denoise Xray fluoroscopy has only been little studied [ALB05, ABL06, DAJW10, CKS93] because it generally requires the transparent nature of Xray images to be taken into account. Motion compensation is more often used in Xray to improve registration of mask and injected images in Digital Subtracted Angiography (DSA) [WSK⁺08, MNV99], or between live fluoroscopy and an anatomical map typically illustrating vessels or heart chambers [BLHS09, BLSH10, Lia10, KBR⁺09, BVG09, AGZ⁺08]. In a more reduced number of studies, motion compensated images sequences are displayed to clinicians with the hope of improving their diagnostic/guidance accuracy [ZAT⁺09, BCD⁺07].

DSE techniques have emerged independently from two different viewpoints. Close et al in [CAW00] were interested in the representation of an X-ray image as a superimposition of transparent layers² and proposed a general framework to estimate the layers present in a sequence of images. They considered that if the motion of a given layer was known, its intensities could be estimated by averaging the value of each pixel along its trajectory in the image sequence. They relied on the assumption that different layers have independent motions. In their processing the tracked layer is integrated constructively thanks to registration whereas the other layers that are not following the same motion tend to be blurred. They proposed an application of their method to stent enhancement, inferring the motion of the layer of the stent by the motion of the balloon marker-balls. In this case, the layer decomposition method boiled down to detecting the successive positions of the marker-balls in the sequence and integrating the images along the resulting trajectory.

Independently and simultaneously, Florent et al [FNLR08] were interested in motion compensated noise reduction and filed a patent describing a method to enhance the visibility of stents. Similarly they proposed to detect and track the balloon marker-balls, to use them to infer the motion of the stent and to average the values of the pixels along their trajectories. They were the first ones to describe an algorithm aimed at detecting the marker-balls with limited user interaction. The general algorithm architecture, common to Close and Florent is illustrated in Fig. 2.3.




FIGURE 2.3: General architecture of a DSE algorithm.

²For more details, the transparent layer formalism is thoroughly described in Chapter 3.

In this chapter, we report on new contributions from both application and technical perspectives. Within the DSE application domain, we note that the following ideas are presented for the first time (i) a complete analysis of DSE from clinical problem statement to performance assessment, (ii) the description of a fully automated algorithm, (iii) an exhaustive validation of this DSE algorithm. Regarding the image processing aspects of DSE, two original contributions are proposed. The first novelty stands in the guide-wire segmentation techniques that are thoroughly described in the chapter. It consists in spanning a set of parametric curves satisfying the physical constraints of our application and integrating profiles along the curve. This integration step is what makes the technique very robust to noise (as demonstrated in the validation section 2.4.1.4). Image noise, as a result of minimizing patient and clinician dose, is a significant obstacle for guide-wire detection. The proposed technique has several nice properties that are demonstrated in the chapter : the segmented guide-wire is smooth, the execution time is reasonable, and it is robust to a large amount of noise. The second novelty with respect to the image processing aspects of DSE is the registration technique. Our analysis of the problem of DSE led us to propose a non-linear registration in a domain where linear registration has previously been considered sufficient. The superiority of our registration is demonstrated on a large number of clinical sequences. This technique fulfills the requirements that make it suitable to register not only stents but also any medical device inserted inside a vessel, or vessels themselves and can thus have various applications.

The chapter is organized as follows: In section 2.3 we define user requirements and flow them down to quantifiable metrics. We show how the main characteristics of a successful DSE algorithm can be deduced by the analysis of the problem. Then we describe a fully automated algorithm to perform DSE that segments the marker-balls and the guide-wire and registers images in a non-linear manner to take into account vessel and stent deformations. In section 2.4, we validate the fulfillment of the user needs on a large number of cases, addressing the performance of the segmentation of the landmarks, the image quality improvements and comparing non-linear to linear registration. In section 2.5, we discuss the results and expose the challenges that are still not met by DSE. Conclusions are drawn in the last section.

2.3 Method

2.3.1 Clinicians' needs

From the clinician's standpoint, the value of a DSE technique depends upon three factors : improvement in image quality, limited user interaction and reasonable execution time. For each factor, we defined a metric and set goals that capture clinical needs.

The first factor, image quality, is the most subjective and therefore the most difficult to measure. For our application we defined a 5-grade image quality scale quantifying the visibility of stents:

1. the stent is hardly visible
2. the stent border can be guessed
3. the stent border is clearly visible
4. the stent border is clearly visible and some struts are visible
5. the stent is perfectly visible

In order to quantify improvement in image quality, two clinicians³ rated images before and after the application of the DSE technique. They considered that a statistically meaningful improvement of one grade is necessary for the method to be adopted. Regarding user interaction, clinical feedback suggested that a DSE technique that is successful in 80% of the cases in a fully automatic manner, and in 90% when the user selects a region of interest (ROI), was sufficient to be regularly employed in practice. In terms of execution time, one has to keep in mind that an angioplasty procedure typically lasts between 20 minutes and 2 hours. During an interventional cardiac procedure, clinicians do not need to have DSE images in real-time, but in a “reasonable” time after the deployment of the stent. The delay induced by the computation of the DSE image should not disturb the regular work-flow of the procedure. Clinical feedback suggests that a DSE technique should not exceed 30s execution time.

2.3.2 From problem statement to algorithm design

The main characteristics of a DSE algorithm that will fulfill the clinicians’ requirements can be deduced from a detailed analysis of the DSE problem. In this section we study how image quality can be improved. In the case of stents, it can be split into four factors : contrast, zoom, sharpness and noise.

2.3.2.1 Contrast

DSE seeks to emphasize the subtle contrast variations induced by the presence of the stent. However, traditional processing and display strategies are generally optimized for the visualization of large structures like coronary vessels while concurrently suppressing noise, and not for the fine stent details. Consequently, contrast and windowing of traditional processing approaches are suboptimal for stent visualization (Fig. 2.4). Dedicated image processing, like unsharp masking (background subtraction), must be set in place

³Dr François Funck, head of the cardiology department of centre hospitalier René Dubos (Pontoise - France) and Neils Guillard, interventional cardiologist in the same institution.




FIGURE 2.4: Left: Illustration of a deployed stent in an x-ray image processed with standard contrast enhancement techniques, which may be considered a form of DSE. The resulting stent visualization is substantially inferior to that produced with a dedicated technique, as shown on the right.

to enhance the fine details of the images.

2.3.2.2 Zoom

The stent only represents a small portion of the original images. Typically a 100^2 area in a 1024^2 image. DSE can directly improve stent visualization by zooming on the region of interest and masking un-necessary image parts. The comparison of Fig. 2.2 (a) and (c) illustrates this.

2.3.2.3 Sharpness

The sharpness of the enhanced image is a function of the sharpness of each input frame and the precision of the registration. The sharpness of the input frames is mainly ruled by the focal spot of the Xray tube and the characteristics of the detector. These physical parameters are generally well optimized by the constructors and the input images can be considered sharp, at the exception of one or two frames every heartbeat that are sometimes blurrier due to heart motion velocity. Sharpness is thus intrinsically related to the precision of the registration. Registration is crucial, considering the stent geometry relative to the optical magnification and detector pixel pitch. Stent struts are typically 0.16 mm wide, less than 0.2 mm pitch [RLM⁺05]. Taking into account a typical magnification factor of 1.4 for cardiac interventions, the projection of the stent struts onto the detector are on the order of one pixel. Finally the stent structure is blurred by the system Mean Transfer Function (MTF) that tends to spread the stent strut over the neighbor pixels. Determining the required precision of the registration and its implications for the design of a stent enhancement algorithm was the subject of dedicated studies. Langan et al [RLM⁺05] determined that DSE is robust to some degree of error in the registration process. They centered their approach on a task oriented framework, evaluating the

ability of an observer to detect stent under deployment. They concluded that it was improved by DSE up to a registration error of 2 pixels but that the benefit exhibited a significant fall off beyond 1.25 pixels of error. On our side, we attempted to quantify the effect of the registration error directly on the image quality regardless of a specific task. We imaged several still stents with an X-ray angiographic system. We averaged these images with various simulated registration error (Fig. 2.5). The registration error was modeled as a uniform distribution on each coordinate of amplitude ranging from 0 to 2.5 pixels. It turned out that stent struts were visible up to an error below 1 pixel. An error between 1 and 2 pixels would preserve the border but make delineation of stent struts very difficult. Any error above 2 pixels would impair the visibility of the stents. These two experiments confirmed that given stent geometry and acquisition resolution, registration errors less than a pixel are tolerable, whereas registration error in excess of 2 pixels severely degrades performance rendering the algorithms benefit to be questionable. We want to emphasize that these results hold for the imaging system considered in our experiment and in the cited study [RLM⁺05]: the Innova 2000 and 2100 series (GE). They shall be revisited for systems with different imaging chains since pixel size and magnification factor can vary. In order to optimize this error we analyzed its origins. We identified two primary root causes for misregistration: landmark localization error; and the error in inferring stent motion relative to that of the landmarks. Stent and landmark motion is complex due to projection of 3D motion onto a 2D plane, and stent motion that is independent of the landmarks (see section 2.5.6). We concluded that, in order to optimize stent visualization, we need to minimize landmark segmentation error (ideally sub-pixel). Additionally, to minimize the stent motion estimation error, we have added the guide-wire supporting the markers balls to the set of landmarks. We show that by considering guide-wire and marker ball motion, we have improved the accuracy and robustness of stent motion estimation (Fig. 2.6).

2.3.2.4 Noise

The major source of noise in X-ray fluoroscopy is quantum noise. It follows a Poisson distribution and it is spatially correlated but temporally white [SBE⁺95, ASS99]. A variance stabilization transform is commonly applied to Xray images [Ans48] after which the noise can be considered stationary and Gaussian. There are two sources of information redundancy in the image sequences that can be used to reduce the noise: the spatial content of each independent image, and the temporal content of the image sequence. The spatial content of a stent image being very complex, intra image denoising could result in stent detail degradation. Therefore a DSE algorithm will preferentially restrict denoising to the temporal domain. The temporal redundancy is advantageously




FIGURE 2.5: Study of the impact of the misregistration error. Image after the application of DSE with various misregistration magnitudes. The images were produced by registering thirty frames with a registration error following a uniform distribution and averaging the frames. The misregistration magnitude is indicated in pixel units below each image.



FIGURE 2.6: Schematic illustration of the limitations of the linear, marker-ball based registration realized with an image processing software for explanatory purpose (not necessarily geometrically exact). The two top images are two images of a same stent. Bottom is the marker-ball based registration. Although the two marker-balls are accurately registered, the stent, that underwent a non-linear deformation is not accurately registered. The guide-wire that supports the marker-balls is also undergoing a non-linear deformation. We use it as additional landmark in our *nonlinear* registration.

The stent image is derived from a photo of a Cypher coronary stent [tim05]

utilized by integrating the values of each pixel of the stent along its trajectory (therefore it requires the estimation of the trajectory of each pixel of the stent). The way the values are combined together and the number of images taken into account drive the noise reduction capability. Typical DSE strategies [RLM⁺05, Koo05, SF09] involve averaging the pixel values over a short (≈ 30 frames) sequence. It yields a noise reduction factor of approximately 5 that is usually enough to make a major difference in image quality.

We study here in more details the way to combine the images and the required number of images.

To the best of our knowledge image combination in the context of DSE has not been studied so far. All the authors [RLM⁺05, KL05, Koo05, SF09] propose to simply average the values of the motion compensated images to produce the DSE image. In the context of Digital Subtracted Angiography (DSA), this topic has been studied in more details [KSW01b, KSW01a, SGH⁺04, LT05]. Authors propose a wide variety of techniques: rank filters [KSW01b, KSW01a], match filters [KSW01b, KSW01a], entropy measurement [SGH⁺04] and learning based approaches [LT05]. However, the problem is totally different from DSE: in DSA authors want to capture the arrival of contrast media that induces a change in the temporal series of values at a given pixel. On the contrary, in DSE the main hypothesis is that there is little change in the motion compensated images object along time. Therefore the image combination techniques we propose in this section are significantly different.

Let us call $\{x_i\}_{1 \leq i \leq n}$ the values of a given pixel along its trajectory in the n images of the sequence. We model the noise on the pixels by considering that the $\{x_i\}_{1 \leq i \leq n}$ are the realization of random variables $\{X_i\}_{1 \leq i \leq n}$. We examine the task of combining the values of the pixels in the different images of the sequences by considering that it consists in estimating a given parameter of the distributions of the $\{X_i\}_{1 \leq i \leq n}$. The choice of the adequate estimator then depends on the hypotheses we make on these distributions. Given the properties of the noise previously listed, we consider the $\{X_i\}_{1 \leq i \leq n}$ to follow Gaussian distributions of same standard deviation σ . Estimators are classically characterized based on their bias and variance. In our case the bias induces a shift in pixel intensities in the combined image and the variance predicts its noise.

We consider two different hypotheses regarding the $\{X_i\}_{1 \leq i \leq n}$ that yield different estimators. The first one is that they all follow the same distribution of mean μ and variance σ^2 . It relies on the assumption that the pixels have been perfectly tracked in the background subtracted images (we call it the hypothesis 1): a same pixel corresponds exactly to a same stent or background element and has thus the same value across the images except for the noise. Under this assumption we aim at estimating the expectation (the parameter μ) of the underlying Gaussian distribution. The empirical mean (Eq 2.1) is the non biased estimator of minimal variance in this case. It is thus an ideal choice. The noise in the combined image⁴ (the standard deviation of the estimator) is $\sigma n^{-1/2}$.

$$\bar{X}_n = \frac{1}{n} \sum_{1 \leq i \leq n} X_i \quad (2.1)$$

⁴This results holds independently of the distribution of the $\{X_i\}_{1 \leq i \leq n}$. It is a consequence of their statistical independence. At this stage the Gaussian hypothesis is thus not useful.

The second hypothesis we consider regarding the distribution of the $\{X_i\}_{1 \leq i \leq n}$ is that they do not all have the same mean. It is an attempt to account for imperfections in the pixel trajectory estimation. These inaccuracies can result from the non uniformity of the cardiac motion. More precisely, in a short lapse of time in the cardiac cycle (typically less than 20%) corresponding to the beginning of the systole, the cardiac motion accelerates and the motion estimation may be less accurate. We thus propose the hypothesis that the values of a pixel along its estimated trajectory correspond to a same mean in 80% of the samples and to different means in the other 20%. In such a case the empirical mean \bar{X}_n may not be the optimal estimator. A common strategy in such a situation where outliers are present in the sample is to use a robust statistic, typically the median and the trimmed mean [Wik11d]. In order to further express these estimators let us consider the rank statistic $X_{(i)}$ that returns the i^{th} value of the set $\{x_i\}_{1 \leq i \leq n}$ sorted in the increasing order. The median is simply $X_{(\lfloor n/2 \rfloor)}$ and the trimmed mean excluding the k lowest and the l highest values on each side is defined following Eq 2.2. In this section the brackets $[\cdot]$ denotes rounding to the closest integer.

$$\bar{T}_{n,l,k} = \frac{1}{n - k - l} \sum_{k+1 \leq i \leq n-l} X_{(i)} \quad (2.2)$$

The trimmed mean is often considered to be in between the mean and the median: for $l = k = 0$ it equals to the mean and for $l = k = \lfloor n/2 \rfloor$ to the median. We selected it for our application since it allows to reject a precise percentage of outliers. We consider three trimmed mean estimators, rejecting $p = 20\%$ of the outliers: the symmetric one $\bar{T}_{n, \lfloor \frac{p}{2}n \rfloor, \lfloor \frac{p}{2}n \rfloor}$, and the two disymmetric ones $\bar{T}_{n, \lfloor pn \rfloor, 0}$ and $\bar{T}_{n, 0, \lfloor pn \rfloor}$. Regarding rank filters, we did not restrict to the median, and for sake of exhaustivity considered the 5 following ones : $X_{(0)}$, $X_{(\lfloor n/4 \rfloor)}$, $X_{(\lfloor n/2 \rfloor)}$, $X_{(\lfloor 3n/4 \rfloor)}$, $X_{(n)}$. We note that $X_{(0)}$, that takes the min of the values of the pixel along the sequence is commonly employed in vascular imaging under the name peak opacification to image contrast media traveling inside vessels [KSW01b, KSW01a].

We thus selected 9 estimators to combine the values of the pixels across the frames : the empirical mean, the 3 trimmed means, and the 5 rank filters. We first compared these estimators based on their bias and variance under hypothesis 1. Since it is delicate to precisely simulate the second hypothesis we applied the estimators to a set of 30 clinical sequences and visually evaluated the image quality of the combined images.

It can be demonstrated that under the assumption of a symmetric distribution (that is the case of the Gaussian) the estimators \bar{X}_n , $X_{(\lfloor n/2 \rfloor)}$, and $\bar{T}_{n, \lfloor \frac{p}{2}n \rfloor, \lfloor \frac{p}{2}n \rfloor}$ are non biased and the other ones biased. Their variance can be analytically derived for \bar{X}_n and $X_{(\lfloor n/2 \rfloor)}$, and the trimmed means respectively $\frac{\sigma^2}{n}$ and $\frac{\pi\sigma^2}{2n}$. However, although the cumulative distributions of the rank filters $F_{X_{(i)}}$ can be derived from the Gaussian cumulated distribution function denoted F according to equation 2.4 [Sv007], the moments are not

straightforward to derive (for instance trimmed means variances are $\frac{\sigma_W^2}{n(1-p)^2}$ where σ_W is a function of the truncated moments of the Gaussian distribution [Sti73]). We estimate them under the assumption of centered normalized and uncorrelated Gaussian variables $\{X_i\}_i$ through simulation (Fig. 2.7).




FIGURE 2.7: Standard deviation of the estimators versus the number of images in the input sequence. Results obtained with one billion simulations. Observe the general decreased in $n^{-1/2}$ predicted by the theory for the empirical mean, the trimmed means and the median.

$$F_{X_{(i)}}(x) = \sum_{j=1}^n n \binom{n}{j} F(x)^j (1 - F(x))^{n-j} \quad (2.3)$$

For \bar{X}_n and $X_{([n/2])}$ the variance follows the expected theoretical curves. The general decrease of the standard deviation in $n^{-1/2}$ is observed for all the estimators except the min and max (resp. $X_{(0)}$ and $X_{(n)}$). We observe that \bar{X}_n has the lowest standard deviation, very close to the one of the trimmed means $\bar{T}_{n, [\frac{p}{2}n], [\frac{p}{2}n]}$, $\bar{T}_{n, [pn], 0}$ and $\bar{T}_{n, 0, [pn]}$. The rank filters have higher standard deviation, and can be ranked in increasing order (denoting $\sigma(E)$ the standard deviation of estimator E):

$$\sigma(X_{([n/2])}) < \sigma(X_{([3n/4])}) = \sigma(X_{([n/4])}) < \sigma(X_{(0)}) = \sigma(X_{(n)}) \quad (2.4)$$

Therefore, we expect to observe an optimal noise reduction capability for the empirical mean, very closely followed by the trimmed means, and larger noise in the rank filter estimates. We validated the relevance of our model on two clinical sequences. We measured the noise in the images by computing the empirical standard deviation in an ROI where there is no significant structures. The background of the images being subtracted, the choice of such an ROI is facilitated (Fig. 2.8). We performed the noise

measurement in one image of the original sequence and on the combined images obtained with each of the 9 estimators (for a combination of 30 images). We then plotted the noise standard deviation predicted by our simulated curves (Fig. 2.7) and the measured ones in Fig. 2.9. The overall accuracy is satisfying: Our model predicts well the general impact of the estimators on noise with an average absolute error of 9%. It validates that the hypotheses on the noise that we made (predominance of quantum noise, Gaussian distribution and temporal independence) are valid enough.




FIGURE 2.8: Application of the background subtraction to a clinical image. Left: Original. Right: Background subtracted. Observe that most of the structures have disappeared.

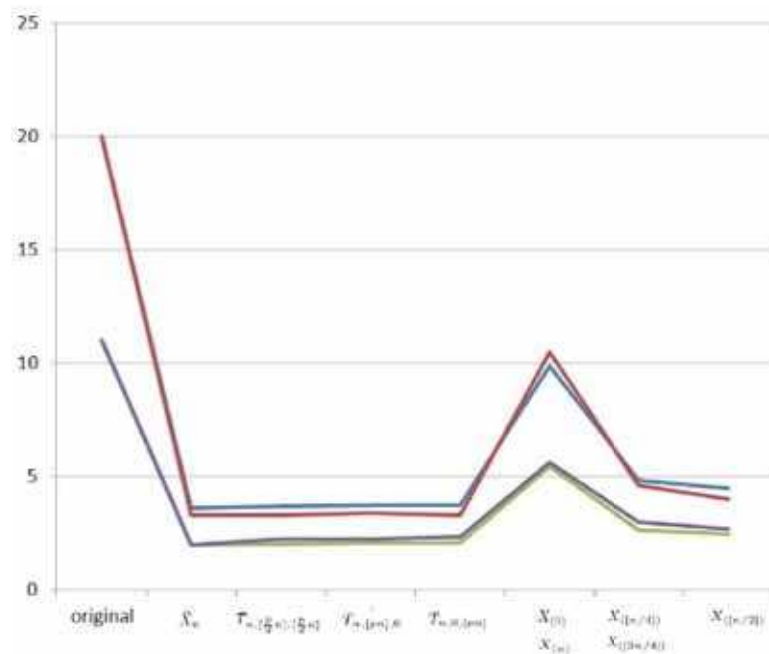


FIGURE 2.9: Simulated vs estimated noise on two clinical cases for the 9 estimators: red and purple are estimated, blue and green are simulated. The overall precision of our model is typically 10%.

We further explored the performance of the different estimators by evaluating the image quality visually on 30 clinical cases. We used \bar{X}_n as reference and performed

relative rating. The observer (who is also the author of this thesis) was asked to rate each estimator among the 5 following categories:

- - - : The image is significantly worse than the one obtained with \bar{X}_n . More details are visible with \bar{X}_n .
- - : The image is slightly worse than the one obtained with \bar{X}_n . The same details are visible but noise or contrast is less good.
- = : No image is better than the other one.
- + : The image is slightly better than the one obtained with \bar{X}_n . The same details are visible but noise or contrast is better.
- ++ : The image is significantly better than the one obtained with \bar{X}_n . Less details are visible with \bar{X}_n .

For this task the examination was on purpose restricted to the stent area. The output of this review are presented in Table 2.1. Since the images were never better than with \bar{X}_n we omitted the empty categories + and ++. The figures demonstrate that the mean filters (empirical and trimmed) perform better than the rank filters. All the mean filters are visually indiscernible in the stent area. Among the rank filters that generally exhibit more noise, the best one is the median that yields results that are never very different from the reference (no case in the category - -). $X_{([n/4])}$ is often acceptable also though significantly worse than \bar{X}_n in some cases. The other rank filters generally yield very poor results. We illustrated the images produced with the different estimators and their difference to \bar{X}_n in Fig. 2.10.

Observation of the part of the image that does not contain the stent (and that is thus less clinically critical) shows that trimmed means are more robust to registration inaccuracies. Far from the markers (typically 50 pixels away from the marker segment) the landmarks are not registered anymore in our experiment and the guide-wire appears at different locations in each image. As a result it appears duplicated in the combined image. Since the guide-wire is darker than the background, it affects the lower part of the distribution of the $\{X_i\}_i$ and trimmed means $\bar{T}_{n, [\frac{p}{2}n], [\frac{p}{2}n]}$, $\bar{T}_{n, [pn], 0}$ allow to attenuate this undesirable effect (and not $\bar{T}_{n, 0, [pn]}$ that keeps unchanged the lowest part of the distribution). We illustrate this observation on a clinical case on Fig. 2.11.

Rating	$X_{(0)}$	$X_{([n/4])}$	$X_{([n/2])}$	$X_{([3n/4])}$	$X_{(n)}$	$\bar{T}_{n, [\frac{p}{2}n], [\frac{p}{2}n]}$	$\bar{T}_{n, [pn], 0}$	$\bar{T}_{n, 0, [pn]}$
- -	100	3	0	27	100	0	0	0
-	0	67	50	67	0	0	0	0
=	0	30	50	7	0	100	100	100

TABLE 2.1: Repartition of the images produced with the different estimators according to relative image quality scale (percentages).

The output of this study on image combination and its effect on noise and image quality is that empirical and trimmed means appear to be the most relevant techniques and yield a noise reduction factor of approximately $n^{-1/2}$ when n images are combined. They are equally relevant on the stent area but the trimmed means are superior in the rest of the image where registration is less accurate. For the remainder of this chapter we will produce results with the empirical mean, since they have been produced before this part of the study was complete. We expect to have optimal results on the most clinically relevant part: the stent.

We have seen that, theoretically, an increase in the number of combined images always decreases the resulting DSE image yielding asymptotically a total noise removal ($\sigma(\bar{X}_n) \xrightarrow{n \rightarrow \infty} 0$). However three factors tend to limit the practical number of images taken into account:

- The improvement in image quality is limited by the minimal image quality required to visualize the stent. Beyond this point, there is no clinically added value to further denoising. Moreover, the inaccuracies of the motion compensation tend to introduce blur into the DSE images that limit the quantity of visible stent details. Those destroyed by this blur cannot appear whatever the noise reduction. Therefore the image tends to converge to an optimal point when the number of images increase.
- The decrease in noise is not a linear function of the number of image combined. The theoretical law of noise decrease being $n^{1/2}$, each additional frame reduces noise less than all the previous ones. With 25 images, 80% of the noise is removed. In order to reach 90% of noise removal one has to combine 100 images.
- Regarding Xray dose, a relevant order of magnitude is the reference level in France for an angioplasty: 1355 cine frames. It is defined as the 3thrd quartile of the distribution, meaning that 75% of the procedures involve less frames than this limit. It is desirable that DSE not increase significantly the whole procedure dose. A few percent of the dose (some tens of frames) are far more acceptable than some tens of percent (some hundreds of frames).

For all these reasons we will only consider short sequences of some tens of frames. In order to evaluate the number of images to take into account in the image combination, we studied the residual noise and the visual image quality as a function of the number of combined images. For this study we used the empirical mean estimator. We studied experimentally the image quality and the noise in 4 short sequences of some tens of frames as a function of the number of combined images. The noise standard deviation measurements, plotted in Fig. 2.12, show that once again the noise follows reasonably

the expected power low. Regarding image quality, DSE images resulting from the combination of 1, 10, 20, 30, 40, 50 and (when available) 60 frames are illustrated in Fig. 2.13. We observe, as expected, that most of the benefit comes from the first tens of frames. Moreover, above 30 to 40 frames the DSE image converges to some optimum. These considerations and results motivate the use of sequences of 30 to 40 frames in DSE. The noise reduction factor is in this case approximately 5.

2.3.2.5 DSE algorithm design take away

To conclude this section, our analysis of the DSE problem resulted in the determination that a successful DSE algorithm shall automatically process short (30 to 40 frames) image sequences with a high success rate ($> 80\%$) and enable a computation time ($< 30s$) to fit within the interventional work-flow. Moreover, in order to attain optimal image quality, it shall perform noise reduction preferentially by temporal techniques, combine images with (trimmed) averaging, include a background removal and adaptive zoom for visualization. Last but not least, it shall detect the marker-balls, the guide-wire and register the stents with a high precision, ideally sub-pixel.




FIGURE 2.10: Result of the different estimators on a clinical case. Left column: result obtained with \bar{X}_n (reference). Center: Result obtained with the considered estimator. Right: difference image. From top to bottom: $X_{(0)}$, $X_{([n/4])}$, $X_{([n/2])}$, $X_{([3n/4])}$, $X_{(n)}$, $\hat{T}_{n, [\frac{p}{2}n], [\frac{p}{2}n]}$, $\hat{T}_{n, [pn], 0}$, $\hat{T}_{n, 0, [pn]}$




FIGURE 2.11: Combined image. From left to right: \bar{X}_n , $\bar{T}_{n, [\frac{n}{2}, \frac{n}{2}]}$, $\bar{T}_{n, [pn], 0}$ and $\bar{T}_{n, 0, [pn]}$. Observe that $\bar{T}_{n, [\frac{n}{2}, \frac{n}{2}]}$ and $\bar{T}_{n, [pn], 0}$ allow to attenuate the undesirable guide-wire multiplication effect far from the landmarks.




FIGURE 2.12: Noise reduction as a function of the number of combined frames (1, 10, 20, 30, 40, 50 and 60 when available) with the empirical mean estimator for four different clinical sequences (one color for each: red, blue, purple and green). Solid line, manual estimation, dashed line theoretical noise with the $\sigma n^{-1/2}$ model. Ordinate: noise standard deviation; Abscissa: number of averaged frames.




FIGURE 2.13: DSE images for different numbers of averaged frames: from top to bottom 1, 10, 20, 30, 40, 50 and (when available) 60. Image have undergone background subtraction as recommended in Section 2.3.2.1. Observe on the very low stent visibility in the original data on top row. Visibility increases rapidly in the first rows, then stagnates.

2.3.3 DSE algorithm

In this section, we describe a DSE algorithm derived from the objectives and constraints previously listed. We first give a general overview of the method, then we detail the parts that are specific to our approach. Finally we explain our strategy to find a relevant setting of the many parameters that are involved.

2.3.3.1 Algorithm overview

The DSE algorithm that we propose can be divided into four main steps (Fig. 2.14): landmark detection and tracking, image registration, image combination, and display processing.

A bottom-up approach is pursued in landmark detection and tracking. First, in each image, we detect points representing potential marker-balls, then we form pairs of points and build tracks of pairs. The detection of the potential marker-balls is close to the problem of detecting micro calcifications in mammograms [CSE06, Gri91]. Similarly, we pre-process images with a dark top hat [Soi99]. It enables the removal of the background variations while retaining dark objects of a given scale. In practice with a relevant setting of the structuring element, the marker-balls, the guide-wire and the stent are preserved (Fig. 2.15). Moreover efficient implementations exist [vH92]. In the resulting image the potential marker-balls can be characterized as being local minima in a given range of intensity. Forming pairs of potential markers and building tracks of pairs along the sequence (Fig. 2.16) is constrained by a priori knowledge of device (stent/marker ball/guide wire) geometry, and the characteristics of cardiac motion (see section 2.3.4). We assign a figure of merit to each track based on its regularity and attenuation properties of its candidate markers. Finally, we identify the most promising track and we segment the guide-wire supporting each pair of markers.

Image registration of the stent relies on a model to infer stent motion from the position of the landmarks. Each image is registered to a reference position in a non-linear fashion. Image combination consists of averaging all the registered frames together to produce a denoised image. Then the contrast is set to enhance stent details. It benefits from the dark top hat pre-processing that removes background variations. Finally zooming is performed centering on the landmarks. The most challenging problems addressed by our algorithm are the segmentation of the guide-wire and the non-linear registration. We first published them in the patent application [VLB06]. We detail them in the following sections.

2.3.3.2 Guide-wire segmentation

Segmenting the guide-wire in an X-ray image is known to be a difficult problem. Indeed, guide-wires are fine elongated structures, typically 3 to 5 pixels wide, of low Contrast




FIGURE 2.14: Block diagram of our DSE technique.




FIGURE 2.15: (a) Original image, (b) top hat transform applied to (a).




FIGURE 2.16: Illustration of the “detection and tracking” brick of our DSE. Top row, from left to right: original image, candidate points, candidate pairs formed out of the points. Bottom row, from left to right: three successive frames and the detected candidate pairs. Among them the tracking was able to build five relevant tracks.

to Noise Ratio (CNR)⁵. A typical CNR is around 3, but may fall below 1 in the most challenging cases. Previous researches on guide-wire segmentation fall into two categories. One category is based upon building a map of the probability that a guide-wire is present [BVG09], and the second category explicitly segments the guide-wire. The second category is of interest for our application. On the one hand, a set of articles [BAG⁺07, HVP10b, HVP10a, HGG⁺09] present sophisticated approaches to guide-wire segmentation based upon a machine learning formalism. They attempted to solve a problem that is more general than ours since they segment the whole guide-wire, whereas we are only interested in a local segmentation at the close vicinity of the marker-balls. On the other hand, Florent [FNLR08] addresses exactly our problem with a minimal cost path approach but does not guarantee the continuity of the first derivative of the guide-wire curve that is required for our application, detailed in section 2.3.3.4. Therefore we developed a new guide-wire detection procedure designed at producing a smooth curve, while being robust to high level of noise and enabling fast computation. Our observation of clinical sequences suggests that the guide-wire between the markers may be modeled as a simple curve, typically a spline. Most of the time it may be approximated by a parabola. This observation is shared in [FNLR08]. In a subset of cases, the guide-wire possesses an inflexion point and may be accurately modeled by a set of two parabolas or a third degree polynomial expression. In the general case, a simple set of smooth parametric curves will describe all the possible guide-wire configurations between the markers, as well as in a short distance beyond the markers. Based on these observations, we decided to specify a discrete set of curves spanning a family of parametric curves (Fig. 2.17). Each curve is piece-wise polynomial. We proceed in two steps in a coarse to fine approach. First the discretization of the family is such that the maximal distance between two curves is 2 pixels. The guide-wire radius being 3 to 5 pixels we are certain that one of the curves is included in the guide-wire. Then we refined the selected curve to sub pixel accuracy with a maximal distance between curves of 0.5 pixels. A figure of merit is assigned to each curve in the set quantifying how well it represents the image content. We retain the parametric curve yielding the highest figure of merit. The challenge of our approach lies in the design of the figure of merit which is robust to the low CNRs of imaged guide-wires. In order to do so, we averaged intensity profiles line segments which are uniformly spaced between the marker balls, perpendicular to the tested curve, and centered on it (Fig. 2.17). The length of the segments shall be long enough to illustrate the guide-wire profile. Therefore, we selected it to be a few times the expected guide-wire diameter. Moreover, in order to capture all the information along the curve, the segments are spaced by one pixel in our experiment. The resulting average profile enables robust contrast estimation due to the reduction in noise. If the curve is

⁵We define the CNR of an object in the following way: Let v be the value of the object, b the one of the background and σ the standard deviation of the noise in the considered area: $CNR = \frac{|v-b|}{\sigma}$.

well aligned with the guide-wire all the profiles are theoretically the same except for the noise. Averaging the profiles reduces this noise and yields a typical guide-wire profile. Conversely, if the tested curve is not properly matching the guide-wire shape every profile is different and averaging does not make any particular pattern emerge from noise. The analysis of the averaged profile is performed in two steps. First, a test is performed on the shape of the profile to determine if it has a significant minimum at its center (Fig. 2.17 (d) and Fig. 2.18). Then contrast is measured by computing the difference between the value at the minimum and the values at the extremities. Evaluations on clinical data have demonstrated that this method enables guide-wire segmentation in extremely noise-corrupted images (Section 2.4.1.4 and Fig. 2.23 (c) and (d)).




FIGURE 2.17: (a) A set of tested parametric curves to segment a guide-wire between two markers. (b) In this clinical image, the dark line along the guide-wire depicts a parametric curve to be tested. The set of dark short segments represent the perpendicular profiles that are computed for this curve. (c) Illustration of the values along one of the profiles of image (b). We see that the noise is high and there is no particular pattern in this curve. (d) Averaged profile. We see that the curve follows the expected pattern: a smooth curve with a significant minimum at its center.

2.3.3.3 Constrained registration derived from a stent deformation model

Once the landmarks (marker-balls and guide-wire) have been segmented, we must then estimate their motion field. Secondly, we must define how to extend this motion field beyond the landmarks encompassing the image area occupied by the stent. Finally, we utilize the motion field for each image to temporally register the image sequence. We denote $(I_t)_{1 \leq t \leq n}$ as the input image sequence of n images and we define I_{t_0} as the




FIGURE 2.18: Top left: an extract from a cine image depicting the guide-wire, the stent and the markerballs. Top right: 3 candidate curves (note that the blue one is well aligned with the guide-wire centerline whereas the red and green ones are not). Center: All the profiles extracted for each curve (same color code as top right). On the abscissa is the signed distance to the curve and on the ordinate the pixel intensity. Bottom: The three average profiles. Only the one of the blue curve that is aligned with the guide-wire exhibits a marked minimum at its center.

reference frame over which all frames are temporally registered. The motion fields are computed with respect to this reference position. The definition of a motion field on the landmarks is straightforward. We force the marker centers to match from one image to another and the guide-wire to match according to the curvilinear abscissa. We define a reference point, O_t , in each image that is the point of the guide-wire between the markers that separates it into two curves of equal length (Fig. 2.19.a). It is equidistant to the two marker-balls on the guide-wire. We compute the signed curvilinear abscissa from this point. We cope with global foreshortening by normalizing the curvilinear abscissa to the total length of the guide-wire between the markers. The motion field registers points relative to their normalized curvilinear abscissa (Fig. 2.19.b).

In order to extend the motion field from the landmarks to the stent we incorporated design constraints reflecting the physical properties of stents. First, stents must be very rigid on their radial axis in order to maintain the desired artery diameter under

the pressure of the artery wall. Consequently, the radius of the stent at a given point must be constant over the image sequence. Second, stents must be flexible on their longitudinal axis in order to bend with the vessels to follow the heart contractions and expansions (Fig. 2.20). Therefore stents follow the deformations of the major axis of the vessel. Finally, we assume that the central line of stents does not shrink when the stents bend (the stents are extending on one side and shrinking on the one, but the length of their centerline is supposed to be invariant). These motion assumptions were confirmed through observation and analysis of clinical image sequences.

Although the major axis of the vessel is unknown, the guide-wire can be used as a surrogate estimate. The first two constraints can be translated into constraints with respect to the position of the guide-wire. Imposing that the distance of any point of the stent to the guide-wire is invariant along the sequence is enough to fulfill the two first constraints. In order to satisfy the stent length invariance constraint, we consider the orthogonal projection of each point of the stent onto the guide-wire (F_P on Fig. 2.19.a) and impose that the curvilinear abscissa of this point be invariant from one image to another. The constrained registration proposed has been demonstrated to perform well in practice (section 2.4.3), although alternate approaches are possible. In the following section, we detail the image registration based on the aforementioned considerations. The design of this image registration technique and its usage are new and particularly tailored to the application of DSE. Moreover, one can observe that the constraints we listed and the solution that we propose in the next section may be extended beyond this context to register other medical tools traveling inside vessels and potentially vessels themselves as summarized in our previous work [BV08b].

2.3.3.4 Non-linear Registration

We build a system of coordinates (Fig. 2.19.a) in each image I_t such that for any point P , d_1 is its distance to the curve defining the guide-wire and F_P the point of the guide-wire that actually minimizes this distance. F_P is unique outside of the medial axis⁶ of the curve. Let us call d_2 the curvilinear abscissa from O_t to F_P along the guide-wire curve divided by the total length of the curve between the markers. We propose to use d_1 and d_2 as coordinate system. In order to create a coordinate system without ambiguities, let us give signs to d_1 and d_2 . Let us call M_1 and M_2 the centers of the two marker-balls in this frame. We define the curvilinear abscissa d_2 in the direction from O_t to M_1 to be positive and the direction from O_t to M_2 to be negative. Noting the curve defining the guide-wire is smooth, we can define a tangent to the curve $\vec{t}(F_P)$ at any point F_P of the curve, pointing in the direction of positive curvilinear abscissa. Finally, we also

⁶The medial axis of an object is the set of all points having more than one closest point on the object's boundary. In 2D, the medial axis of a plane curve S is the locus of the centers of circles that are tangent to curve S in two or more points, where all such circles are contained in S [B⁺67].

define a normal vector $\vec{n}(F_P)$ such that $(\vec{t}(F_P), \vec{n}(F_P))$ is direct and we assign d_1 the sign of the scalar product $\vec{t}(F_P) \cdot F_P \vec{P}$. Outside of the medial axis of the curve defining the guide-wire, T_t , that maps a point P of I_t to (d_1, d_2) , is a continuous one to one mapping but is not defined on the medial axis. On a database of 241 clinical cases, we observed that the stent is never intersecting the medial axis. This can be explained by the rather high radii of curvature of the guide-wires in the clinical images. Therefore there is no artifact introduced by this transformation on clinical images in the area of interest.

In order to perform registration, we need to define the inverse transform T_t^{-1} that maps (d_1, d_2) to a point P . The definition of T_t^{-1} is straightforward. From the point O_t , we define the point F_P as the only point on the guide-wire having the signed curvilinear abscissa d_2 . Then we define the point $P = T_t^{-1}(d_1, d_2)$ by moving away from F_P perpendicularly to $\vec{t}(F_P)$ by (signed distance) d_1 . The registration of a point P of the image I_{t_0} on an image I_t is the point P'_t computed according to:

$$P'_t = T_t^{-1}(T_{t_0}(P)) \quad (2.5)$$

Conversely any point P'_t of an image I_t is registered to the point P of I_{t_0} according to:

$$P = T_{t_0}^{-1}(T_t(P'_t)) \quad (2.6)$$

Let \tilde{I}_t be the result of applying equation (2) to every point in I_t . In any image \tilde{I}_t , the marker-balls and the guide-wire are exactly at the same location as in I_{t_0} . Using the empirical mean estimator, the image combination simply consists in averaging all the images \tilde{I}_t to produce the enhanced image $I_{enhanced}$:

$$I_{enhanced} = \frac{1}{n} \sum_{t=1}^n \tilde{I}_t \quad (2.7)$$




FIGURE 2.19: (a) Notations for the curve based registration. (b) Schematic illustration of the motion field on the landmarks



FIGURE 2.20: Coronary stents are very rigid on their radial axis to handle the pressure of the artery wall, but very flexible in the perpendicular direction to conform to the curvature of the artery.

2.3.4 Settings and algorithm validation

Detecting the marker-balls, forming pairs, tracking them, and segmenting the guide-wire required setting a large number of parameters. In order to determine optimal parameter settings, we collected a large database of 241 clinical sequences, representing approximately 7000 marker pairs. An operator manually determined the position of each marker ball in each image, and marked some guide-wire points in a subset of the images. Over this database we have been able to estimate typical values and variability of marker characteristics, marker motion, and guide-wire curvature. Moreover it enabled testing the landmark detection performances for a given set of algorithm parameters by comparing the results given by the algorithm to the ground truth. This process enabled parameter tuning for the detailed DSE algorithm. We found that the optimal setting of some key parameters is as follows: the size of the structuring element for the subtraction must be set to 9 pixels, the maximum height of the parabola describing the guide-wire must be set to 20% of the distance between the markers, and the distance between two markers does not vary more than 15% from one frame to another. The results presented in the next section have been obtained with the best parameter set that we tested.

2.4 Results

According to the clinical criteria defined in section 2.3.1, we elaborate on the validation framework for our algorithm. We have quantified its ability to automatically segment the landmarks. Results regarding the precision and robustness of their segmentation are also reported. Clinical interest of the technique has been assessed by evaluating the improvement it brings in image quality. In order to compare to other existing methods we have assessed the performance of the linear marker-ball based registration versus our non-linear registration and quantified both the difference in motion fields that are encountered on clinical data and difference in image quality in the resulting images. Finally we compared our software to another commercially available solution.

2.4.1 Landmark detection and localization performance

2.4.1.1 Accuracy and robustness evaluation strategy

We have separated the problem of estimating the landmark detection/tracking performance into two sub-problems. The first category is to quantify, on clinical sequences, how well the automatic marker-ball segmentation is able to detect and track the marker-balls. The main challenges are robustness to complex image content (anatomical structures and medical tools) in the detection algorithm, and to cardiac and respiratory motion with respect to tracking. It is evaluated in Section 2.4.1.2. The second category addresses quantitative accuracy of the segmentation, given that the marker balls are detected by the algorithm. We noted in the section 2.3.2.3 that we would like to reach a sub-pixel precision to insure the sharpness of the DSE image. Quantitative performance measurement of the registration accuracy requires precise knowledge of (sub-pixel) ground truth. This precision cannot be achieved by a human operator on clinical sequences. Therefore we set up a specific simulation to assess that is the subject of the section 2.4.1.3.

2.4.1.2 Marker ball detection robustness

In the present section, we deal with landmark detection and tracking performance. To this end, we had an operator mark the position of the markers on clinical image sequences. The position marked by the operator lies inside of the marker-ball but is not precisely at its center since this exact position is unknown. We compared the position of the detected markers to this ground truth on every image, calculating the number of false negatives, true positives and false positives (there are no true negatives since the marker-balls are present in every frame of the sequence) according to the following definitions:

- True positives are frames where the detected marker candidates are *closer* to the ground truth than a given tolerance.
- False positives are frames where the detected markers are *further* from the ground truth markers than a given tolerance.
- False negatives are frames where no marker is detected.

Marker ball being ellipses of large axis between 7 to 10 pixels, the tolerance was set to ± 5 pixels. We computed the number of cases where the markers that our DSE algorithm output were closer than the tolerance to the ground truth. On a database of 241 clinical sequences, illustrating 7230 pairs of markers, our algorithm exhibited a true positive rate of 91.5%, a false positive rate of 1.5% and a false negative rate of 7.5%. We studied the detection performance in more detail to tell if some frames are regularly failed in every image sequence, or if some images sequences completely failed whereas others

were completely successful. The histogram of the true positives, false positives and false negatives per image sequence are reported in Fig. 2.21. We can observe that some frames are regularly missed in each sequence as false negatives (2 to 3 frames). This does not have a strong impact on the final image quality. They often coincide with the small number of frames where the markers are very blurry due to the heart acceleration at the end of the rest phase. These images contain only very low information for stent enhancement, since the stent is also blurred. Regarding false positives, we can notice that in 98% of the sequence there are less than 5% of them. The overall performance over the database is very satisfying, and far beyond the 80% acceptance criterion set in 2.3.1 on the true positive rate.

2.4.1.3 Marker ball segmentation accuracy

For this task, we simulated a marker ball on a guide-wire using a model of a sphere and of a cylinder corrupted by typical noise and system point spread function (see Fig. 2.22). We set the CNR of the synthetic marker-balls to values ranging from 5 to 10, since it is the range of values observed in clinical sequences. We applied our technique to the synthetic images and estimated for each input CNR the average distance between the detected marker location and the real marker center. A marker was considered detected if its center lied within the synthetic marker ball. The marker detection rate was 98%. This excellent result can be explained by the simplicity of our synthetic sequences. A more meaningful success rate is the one of 91.5% reported previously reported in Section 2.4.1 computed on the clinical database. The synthetic sequences are however particularly interesting to estimate the precision of the marker detection. We found that the average error is less than 0.5 pixels both along the guide-wire and perpendicularly to it. As one can expect from the structure of the test images, the tests demonstrated that our marker localization process is more precise perpendicular to the guide-wire than parallel to it.

2.4.1.4 Guide-wire segmentation accuracy and robustness

Two main factors impact the success of guide-wire segmentation. On the one hand, the precise shape of the guide-wire, since we assume that it can be modeled by a set of parametric curves. On the other hand, the contrast to noise ratio of the guide-wire. In order to address the topic of the shape of the guide-wire, we have selected 10 clinical sequences out of 241 on the criteria of very diversified guide-wire shapes and orientations, and rather favorable image quality. We ran our segmentation algorithm on each of these sequences and visually assessed the quality of the result. All segmentations were visually perfect (Fig. 2.23 (a) and (b)). This experiment validated that the set of selected parametric curves spans the natural variability of guide-wire configurations. We




FIGURE 2.21: In dark (top), histogram of the percentage of the true positive per image sequence. In gray, histogram of false positive per image sequence. In white (bottom), histogram of false negative per image sequence.

then considered these segmentations to be the ground truth for these images. Given the guide-wire width of 3 to 5 pixels and the difficulty to delineate precisely its centerline, we considered that this segmentation was precise within a range of plus or minus 1.5 pixel. Recognizing the challenge in the segmentation of the guide-wire being attributed to the low CNRs often encountered in clinical conditions, we degraded the 10 clinical




FIGURE 2.22: An image of a synthetic marker-ball and guide-wire used to quantify the precision of the marker-ball segmentation.

images with noise to evaluate the robustness of the segmentation (Fig. 2.23 (c) and (d)). We ran our algorithm on each degraded image and computed the Hausdorff distance of the segmentation to the ground truth (Fig. 2.23 (e) and (f)) according to Eq 2.8. We generated 1400 degraded images adding noise of various standard deviations, up to producing images with 3 times the original noise level, which is considered an extreme situation. The CNR of the guide-wires ranged from 0.6 to 5.6 with an average of 1.7 ± 0.9 . Over all the degraded images the Hausdorff distance was 0.8 ± 1.2 and was below 1.5 pixels in 93% of the cases. This error increased with an increase in noise. For the cases where the noise was less than twice the original noise its repartition was 0.5 ± 1.0 , below 1.5 in 97% of the cases. When the noise was two to three times the original noise it was 0.9 ± 1.3 , again below 1.5 in 91% of the cases. An analysis of the histogram of the Hausdorff distances demonstrates that the overall 6% of failed segmentations fall equally into two categories (Fig. 2.24). Half of them are the tail of a distribution centered on the true segmentation. The other half are very large errors (superior to 7 pixels) where the segmentation departed from the guide-wire. In these cases, the algorithm was influenced by the image content, for example the border of a stent. The overall results are very satisfying with the segmentation being accurate in 97% of the cases portraying realistic conditions, and in 91% in extremely difficult cases which are less commonly encountered.

$$d_H(X, Y) = \max(\sup_{x \in X} \inf_{y \in Y} d(x, y), \sup_{y \in Y} \inf_{x \in X} d(x, y)) \quad (2.8)$$

2.4.2 Image quality improvement

The image quality has been rated independently by two clinicians, before and after the application of the stent enhancement software according to the 5 grade image quality scale presented in section 2.3.1. A comparison performed on 80 clinical sequences was the object of a previous publication [FGBV09]. A paired T-test demonstrated that an improvement of 1.0 point out of 5 is statistically significant ($p \leq 0.001$). A larger study including 100 patients and 196 sequences confirmed that our DSE technique improves




FIGURE 2.23: Figures illustrating the characterization of the guide-wire segmentation. (a) Input clinical image, (b) Guide wire segmented on image (a), (c) Image (a) degraded up to three times the original noise, (d) Guide wire segmented on image (c), (e) Another input clinical image, (f) Ground truth (in black) and actual segmentation (in white) of the guide-wire in image (e) degraded. The Hausdorff distance, of 12 pixels in this case, is illustrated by the dashed white line.

image quality by more than one point ($p \leq 0.001$) on the same image quality scale over a large variety of imaging conditions and stents. We present here this larger study into more details.

2.4.2.1 Data acquisition

The image sequences included in this study come from patients who were having a stent implanted during an angioplasty procedure on an Innova 2100^{IQ} (General Electric, Milwaukee, WI, USA) at Hospital René Dubos (Pontoise, France). The data collection




FIGURE 2.24: Histogram of the Hausdorff distances in pixels between the ground truth guide wires and the ones segmented on degraded images.

has been performed in the context of an observational study (non-interventional study⁷), following the standard diagnostic and treatment practices. The anonymous images were processed with our algorithm embedded in StentViz (a DSE software from GE) outside of the care workflow. The study included any image sequence of approximately 30 frames, with the deflated balloon in place inside a deployed stent.

From an image quality standpoint, the Innova system was set so as to produce images at 15 fps in the low dose setting. This setting is typical of a cardiac site, and for a DSE software it is a rather difficult case since frame-to-frame marker-ball motion is higher at 15 fps than at 30 fps and image quality is more challenging in the “low” than in the “normal” dose setting (the dose is reduced by a factor of 2). The patient population that was included in this study is typical of the clinical site: the average cardiac patient is male, 60 ± 10 years old, approximately 25% of them being classified obese.

For most of the patients, several sequences were collected, because fluoroscopy and cine acquisitions were performed and/or because several stents were placed. The set of sequences includes 110 fluoroscopy and 86 cine. In 72 cases, a cine and a fluoroscopy of the same patient at the same stage of the intervention have been collected (representing a total of 144 sequences). It enabled comparing the performances of DSE in both acquisition modes.

Stents data, including type, length and diameter, have been reported in 85% of the sequences. The lengths of the stents range from 8 to 34 mm, with an average stent length of $19.4 \text{ mm} \pm 6.9$. The stent diameters range from 2.25 to 4.5 mm, with an

⁷This terminology does not mean that we are not dealing with interventional cardiology, but means that the patients are treated in the standard way.

average diameter of 2.9 mm \pm 0.47. The types of stents imaged in the study are reported in Table 2.2. The two most represented stents are Cypher (44 times) and Taxus (41 times). Several sequences can contain images of the same stent, in the same patient, but under different imaging conditions (angulation or Xray dose).

Stent	Manufacturer	Number of cases	Input IQ	DSE IQ	Δ IQ
Cypher	Cordis	44	1.6	3.4	1.8
Taxus	Boston Scientific	41	1.2	2.4	1.3
Pro-kinetic	Biotronik	23	1.2	2.4	1.3
Vision and Xience	Abbott Vascular	25	1.4	3.2	1.8
Driver and Endeavor	Medtronic	27	1.7	3.6	1.9
Lekton Motion	Biotronik	4	1.0	2.5	1.5
Flexmaster	Abbott	2	1.0	3.0	2.0
Tsunami	Terumo	1	1.0	3.0	2.0

TABLE 2.2: Stents in the image quality study: stent model, manufacturer, number of occurrences, input image quality, image quality after the application of DSE and improvement in image quality.

2.4.2.2 Image quality and clinical benefit results

Each observer was asked independently to review the original image sequences and the resulting StentViz image and to rate stent visibility in both according to the IQ scale defined in Section 2.3.1. The IQ over the whole dataset was 1.2 ± 0.4 before application of DSE and 2.9 ± 1.1 after (Fig. 2.25). The average improvement is 1.8 ± 1.1 ($p \leq 0.001$). According to the IQ scale, DSE allowed in average, improving from an image where the stent is hardly visible to an image where the stent border is clearly visible. A 95% Confidence Interval (CI) for the improvement is [1.7, 1.9]. We assessed inter-observer variability by computing linearly weighted kappa (κ) [VG05]. Since the unprocessed and processed images are of different nature, we computed it separately on both. The κ measuring agreement between the two observers at rating unprocessed images is 0.34 demonstrating “fair agreement”. On processed images the κ of 0.49 is the expression of “moderate agreement”. Finally, the Pearson correlation coefficient between the scores before and after application of DSE is 0.2 ($p \leq 0.001$). This positive correlation demonstrates the general trend: the higher the image quality of the input image sequence, the higher the image quality of the DSE image.

Fluoroscopic sequences included in the study present the advantage of requiring less Xray dose but the drawback of having lower image quality. We investigated how it interacts with DSE. We have compared the performance of DSE on fluoro versus cine by comparing the IQ scores of the 72 cases where both have been acquired in comparable conditions. The fluoroscopic sequences improved from 1.3 ± 0.6 to 2.7 ± 1.1 after processing and the cine from 1.5 ± 0.7 to 3.2 ± 1.0 . We have confirmed that these differences are statistically meaningful by performing paired T-tests. The input cine sequences have

better IQ than the input fluoroscopy ($p \leq 0.001$, 95% CI [0.2, 0.4]). As expected, the DSE images derived from cine sequences have better IQ than the DSE images derived from fluoroscopic sequences ($p \leq 0.001$, 95% CI [0.3, 0.6]). The histograms of the distributions are represented on Fig. 2.26. We observe a shift of population towards higher score passing from fluoroscopy to cine. In order to quantify the proportion of sequences where cine is better than fluoroscopy, we computed the histogram of the difference of the IQ score in cine and in fluoroscopy averaged over the two raters (Fig. 2.27). In 72% of the cases the difference is between -0.5 and $+0.5$, corresponding to very similar results. In 28% of the cases the difference is greater or equal to 1 meaning that the DSE on cine is better than on fluoroscopy.

In order to go beyond pure IQ rating and evaluate the clinical benefit, we asked the observers to rely on their experience, and to assess it based on the couples of processed and unprocessed images. Indeed DSE was not available at the time of the interventions and all the study has been performed in post processing. Each observer answered the yes/no question: “Would StentViz have helped you perform a better angioplasty?”. The word “better” was to be understood as: faster, with more confidence or with less dose. In 13% of the cases, both observers showed agreement and answered *Yes*. Observer 1 alone answered positively in 40% of the cases and observer 2 in 19%. A possible explanation for this difference is that observer 1 was more familiar with DSE at the time of the study. Indeed, he was also working, on a regular basis, with a product from Philips that presented a DSE feature. This may be a reason why he gave more credit to DSE images and therefore found them more useful.

Finally we investigated how the stent type impacts the result of DSE. It is of daily clinical experience that different stents have different aspect under X-rays. Indeed, stent composition and strut geometry impact directly their visibility. We have computed separately the IQ improvement for each stent type (Table 2.2). The stents that are based on the same platform have been grouped together (Driver and Endeavor on the one side, and Vision and Xience on the other one). The results confirmed that all stents are not equally visible in the input images. Moreover the study demonstrates that the improvement is not constant for each stent but depends on its type. In general, we observed that the better the input IQ, the better the output IQ and the larger the improvement. The stents for which we have sufficient data points could be ranked according to their visibility by decreasing order: Driver and Endeavor, Cypher, Vision and Xience, Taxus, Prokinetic.




FIGURE 2.25: Image quality scores histograms. In gray the values before DSE and in black after.




FIGURE 2.26: Image quality score histograms in fluoro and cine separately before and after DSE (StentViz).

2.4.3 Non linear versus linear registration

In this section we studied the impact of the choice of the registration technique (marker based or curve based) on DSE published in [BV08b]. On the first hand we computed the difference it created on the motion fields. On the second hand we studied its impact on image quality thank to image reviews.

2.4.3.1 Quantitative comparison of the transforms

We quantified here the differences between the similarity registration based on the markers and the proposed curve-based registration taking into account the markers and the guide. The increased level of complexity brought by the curve-based transform only makes sense if it differs significantly from the similarity on a relevant proportion of clinical cases. One notices that the 2 proposed transforms are obviously equal on the




FIGURE 2.27: Histogram of the DSE score in cine minus the score in fluoro averaged over the two raters.

marker-balls, since both of them register the marker-balls. It is then natural to wonder how much and where the transforms differ. Both transforms being smooth one can figure out that if the markers are close one to another (in the case of a short stent) the 2 transforms that are equal on the markers do not differ much on the whole stent. In this section we aim to automatically and quantitatively compare the transforms.

The only way to measure this is to have realistic motion and deformation of the stent, the guide-wire and the markers. We have used to this end a database of 144 clinical sequences of deployed stents with the balloons markers and guide-wires, that is a subset of the database presented in 2.4.2.1. We applied our marker and guide-wire detection algorithm to every sequence. In each sequence we chose a reference frame. We then registered every image of each sequence to the reference frame with the similarity transform and with the curve-based transform. Finally we averaged the registered images for each transform to produce the enhanced images.

For comparison purpose we denote $S_{t_0 \rightarrow t}$ the similarity transform that is used to register the frames in [CAW00, FNLR08, Koo05] and that actually registers the marker-balls. It is the composition of a translation that registers the center of the markers, a rotation that aligns the axes of the markers and a scaling that actually registers the markers. Let us denote P_t'' the point of I_t resulting of the application of $S_{t_0 \rightarrow t}$ to the point P of I : $P_t'' = S_{t_0 \rightarrow t}(P)$.

For a given pixel to register, each transform generates a motion field that can be represented by a vector. We compared these vectors at any point P of I_{t_0} by computing the norm of the difference between the vector obtained with the similarity transform, $P\vec{P}_t''$, and with the curve-based transform, $P\vec{P}_t'$. In order to quantify the difference for a whole sequence we averaged $\|P\vec{P}_t' - P\vec{P}_t''\|$, along the sequence. We call the resulting image the transform difference map: I_Δ . Let $I_\Delta(P)$ be its value at a pixel P , it is defined

by the equation:

$$I_{\Delta}(P) = \frac{1}{n} \sum_{t=1}^n \left\| P\vec{P}'_t - P\vec{P}''_t \right\| \quad (2.9)$$

An example of I_{Δ} can be observed in Fig. 2.28 (b). Black represents 0 and white any value above 2.75 pixel. The two dark halos demonstrate that the two transforms are equal around the markers. We can observe that the difference increases further away from the markers, especially in the area of the stent. In order to build a quantitative metric of the difference between the two transformations, we then defined a Region Of Interest (ROI) where we compare both transforms. This ROI must exclude the markers where there is obviously no difference and must take into account the area where we suppose the stent is. We chose a band around the guide-wire, 30 pixel wide, centered between the markers and whose length is half the length of the guide-wire between the markers. Fig. 2.28 (c) depicts such an ROI on a clinical example. Our quantitative measure to compare the 2 transforms is the average in this ROI of I_{Δ} . We denote this quantity Δ . Let $|ROI|$ be the number of pixels in the ROI, Δ is computed by:

$$\Delta = \frac{1}{|ROI|} \sum_{P \in ROI} I_{\Delta}(P) \quad (2.10)$$

Fig. 2.29 summarizes the measures of Δ over the database. In 2.3.2.3 the precision of the registration versus typical stent dimensions is discussed. We found out that although stent struts are typically below the pixel, the observer who is interested in the overall shape of the stent can cope with a registration error up to the value of 2 pixels. Meanwhile the blurring introduced by mis-registration impairs the visibility of the struts of the stent even with mis-registration errors of approximately one pixel. The histogram of Δ demonstrates it is above 0.75 pixels for 68% of the sequences of the database. In 16% of the sequences Δ is superior to 1.5. The two transforms are thus clearly different in a wide proportion of clinical sequences. On a significant number of them we expect great differences in the enhanced images. On Fig. 2.28 (d) for instance, that depicts the guide-wires of all the images of the sequence registered with the linear transformation, we can observe a misregistration of several pixels. By construction these guide-wires would perfectly superimpose with the our linear registration.

2.4.3.2 Impact of the transform on image quality

In this section we evaluated the impact of the choice of the registration on the enhanced image. We visually inspected the enhanced images with both registration techniques on a database of 144 sequences. Each sequence contains 30 frames. We reported if there was a significant difference on the stent and if the stent was better visualized with one of the registration techniques. The observed differences fell into two categories.




FIGURE 2.28: (a) Image obtained with StentViz; (b) I_{Δ} ; (c) ROI to compute Δ . Observe that on the markers the two motion fields are identical (black spots in (b)) but differ on the rest of the image. (d) All the segmented guide wires registered by the marker ball based techniques. Note that they are not accurately registered outside of the markerballs.





FIGURE 2.29: Histogram of the measures of Δ over the database, indicating the average difference between the transformations.

Category A regroups the cases where the same details are visible but where one technique creates more blurring on the stent. Category B regroups the cases when one technique is so much better suited than the other one that we actually see more details with it.


Table 2.3 summarizes the results and tends to demonstrate that the proposed non-linear curve-based technique performs significantly better than the linear marker-ball based registration. Indeed for 27.9% of the cases the non-linear registration produces a better image than the linear one, whereas the opposite only happens in 3.4% of the cases. Moreover in 11.6% of the cases the non-linear registration is better and differs from the linear registration with a difference of category B. The occurrence rate of the opposite is less than 1%. Fig. 2.30 depicts examples where the non-linear registration is better suited than the linear one. Around the markers the two enhanced images are very similar, but on the center of the stent the linear registration produces a strong blurring and details are lost. On the contrary, these details are enhanced with the curve-based registration.

TABLE 2.3: Influence of the transform


curve-based better		=	similarity (marker based) better	
B	A		A	B
16.3%	11.6%	68.0%	2.7%	0.7%




(a)



(b)



(c)



(d)

FIGURE 2.30: Comparison of the linear and the non-linear registration on two sequences where the difference is of category B. (a) Sequence 1 processed with the linear registration, (b) Sequence 1 processed with the *non* linear registration, (c) Sequence 2 processed with the linear registration, (d) Sequence 2 processed with the *non* linear registration.

2.4.4 Execution time

Our DSE technique implemented on a PC with a 3.2 GHz single-core processor performs on average in 16.6 s over a database of 40 cases with a standard deviation of 3.0 s. The maximum execution time was 30s and the minimum was 13s. The algorithm is fast enough to fit seamlessly in the interventional work flow. With a dedicated implementation, on a multi-core Xeon X3450 processor with a 2.66 GHz clock speed, the average time drops down to approximately 3 s.

2.4.5 Comparison with existing software

We led a study to compare the output of StentViz with the one of another commercially available DSE software in clinical conditions with our partners from Hopital René Dubos. This study was originally conducted to support the submission of StentViz to the FDA [Mor09]. In this context, in order to introduce StentViz on the American market, we had to demonstrate its “substantial equivalence” to an existing device, called “predicate device” in the FDA terminology and we chose to compare to StentOp from Paieon. In order to demonstrate equivalence, we gathered a set of images and processed them with StentViz and StentOp. Then we organized blinded reviews with three clinicians independently where they rated and compared the image quality of the images produced by both softwares. The conclusion of the test was grounded on statistics and hypothesis testing.

2.4.5.1 Data collected and stent information

The study included 41 image sequences, originating from 28 different patients. They have been acquired on an Innova 2100^{IQ} system. Some of the most commonly used coronary stents worldwide have been imaged in the study. The following tables show the types of stents, the number of occurrences and the stent dimensions (table 2.4 and 2.5). They are typical of the classical distribution of stent types and dimensions encountered in common clinical practice.

TABLE 2.4: Stents types in the study

Stent Model	Manufacturer	Occurrence in the study
Prokinetic	Biotronik	16
Xience	Abbott	9
Cypher	Cordis	3
Tsunami	Terumo	5
Chrono	Sorin	3
Taxus	Boston Scientific	5

TABLE 2.5: Stents dimensions in the study

	Minimum	Maximum	Mean	Standard deviation
Stent diameter (mm)	2.25	4.00	2.85	0.44
Stent length (mm)	8.00	31.00	16.17	6.57

2.4.5.2 Comparison method

The same clinical sequences were processed with StentViz and StentOp. The observers (Reader 1: Dr Funck from René Dubos Hospital, Pontoise, France, and two cardiologists of his team, reader 2 and 3: Dr Guillard and Dr Decalf) rated independently, and blindly the visibility of stents in the input cardiac record sequences, in the StentViz and in the StentOp images according to the image quality scale we presented in 2.3.1 that ranges from 1 (stent hardly visible) to 5 (perfect stent visibility).

The observers were also requested to compare the visibility of the stents with StentViz and with StentOp, in a side-by-side blinded review and to rate it according to the 5 following classes:

- “Left ++”: The image on the left is better than the one on the right because it enables to see *more details*.
- “Left +”: The image on the left is better than the one on the right but *the same details are visible* on both images.
- “=”: Both are equal.
- “Right+”: The image on the right is better than the one on the left but *the same details are visible* on both images.
- “Right++”: The image on the right is better than the one on the left because it enables to see *more details*.

2.4.5.3 Results of the comparison

Table 2.6 reports the average image quality scores of the input sequences and of the StentViz and StentOp images. μ denotes the average image quality score and σ its standard deviation.

TABLE 2.6: Image quality ratings

	Reader 1		Reader 2		Reader 3	
	μ	σ	μ	σ	μ	σ
Innova	1.39	0.54	1.85	0.73	1.80	0.68
StentOp	1.98	0.76	2.46	0.78	3.05	1.00
StentViz	2.63	0.80	3.39	0.89	3.56	1.07

To demonstrate the substantial equivalence in image quality of StentViz and StentOp, we ran a paired T-Test with the following hypotheses:

- H_0 : average of StentViz scores = average of StentOp scores.
- H_1 : average of StentViz scores > average of StentOp scores.

The hypothesis H_0 has been rejected by the T-Test (p value < 0.001) for each observer. Figure 2.31 depicts the histogram of the difference of image quality scores (StentViz score – StentOp score). H_0 is represented by the blue dot and a 95% confidence interval for the difference of the scores is represented in red. According to this result we can state that in average the StentViz scores are superior to the StentOp scores and that this difference is statistically significant.




FIGURE 2.31: Histogram of StentViz image score - StentOp Image score; Hypothesis H_0 is represented by the blue dot. 95% confidence interval for the difference of scores is represented in red (rated by: Left reader 1, center reader 2 and right reader 3).

The results of the side-by-side comparison are presented in table 2.7.

TABLE 2.7: Comparative ratings

	StentViz ++	StentViz +	=	StentOp +	StentOp ++
Reader 1	10%	61%	22%	7%	0%
Reader 2	27%	61%	12%	0%	0%
Reader 3	24%	39%	34%	2%	0%

We can observe that the results are very favorable to Stentviz, that is perceived superior to StentOp in 74% of the cases. Conversely, StentOp was considered superior to StentViz only in 3% of the cases, and both were considered equal in 23%. Fig. 2.32 presents some images to illustrate the results of the two DSE products.

2.5 Discussion

2.5.1 Non linear versus linear registration

Although section 2.4.3 demonstrated that the non-linear registration outperformed performed linear registration, we must keep in mind that it is more computationally intensive. Indeed, the non-linear registration requires precise segmentation of the guide-wire to compute the motion fields using curvilinear abscissa and projections of points onto




FIGURE 2.32: Two examples ((a) and (b)) and the images produced by StentViz ((c) and (d)) and StentOp ((e) and (f)). Note that with StentViz the details of the stent are more visible.

curves. Nevertheless, it is noteworthy that the overall execution time of the non-linear registration DSE is compliant with the clinicians needs. However some users may prefer an alternate trade-off between image quality and latency within the application.

An additional performance comparison between marker-ball based and guide-wire based registration may be made with respect to the accuracy of the stent diameter. Since clinicians can perform measurements on the enhanced stent to assess if it is properly

deployed [Koo05, CAC⁺09], a registration that preserves the actual diameter of the stent is preferred. As far as the stent follows the motion model described in 2.3.3.3, the guide-wire based registration preserves the stent diameter. On the contrary, there is no such guarantee with the marker ball based registration. Indeed, the linear transform includes a scaling to cope with the change of distance between the markers from frame to frame. The change in distance between the marker-balls can be the result of a change of curvature of the guide-wire due to the heart beat or a change of the angle between the markers and the image plane. In the first case, scaling the whole region of interest is not adequate because it will modify the diameter of the stent. Guide-wire based registration enables better image quality and probably more accurate stent diameters, whereas marker ball based registration results in faster processing.

2.5.2 Comparison with related work

2.5.2.1 Comparison of the technical aspects

There exists in the literature four teams that described similar techniques we can compare to. The first one, [CAW00, CAW02, CWA03], deals with transparent layer decomposition. These pioneer articles do not aim at presenting an automatic DSE algorithm. They demonstrated that motion compensated temporal denoising based on the motion of the balloon can improve the visibility of stents and help a user to position more accurately its edges. The second team published a patent [FNL08] describing a method to perform motion compensated noise reduction based on the detection of one or several markers. They work for the firm, Philips, that commercializes the DSE software called StentBoost. They cite the application to a pair of marker-balls for stent enhancement and describe a semi automatic method (it requires the user to input a region of interest). It is similar to the one described here from a structural standpoint: marker candidates are detected, then pairs are formed in each image. The main difference of their approach versus ours is that they choose the best pair of markers independently on each frame without forcing the pairs to have a smooth motion within the image sequence. We notice that they pay particular attention to designing a marker ball detector that will not produce responses along uniform ridges and that they segment the guide-wire between the markers. Additionally, although they segment the guide-wire, they do not use it for registration. Because the patent only provides a description and no result or evaluation of the technique, comparison is limited. In a separate patent [MF06] dedicated to guide-wire tip tracking and enhancing they describe the segmentation of a guide-wire and its use for image registration. However, they do not specify the registration technique nor present result. They also published on the use of marker-ball detection and tracking for 3D stent reconstruction and compared it to IVUS [SFL⁺09b]. In this publication they briefly mention a marker-ball detection and tracking algorithm that relies on the

temporal smoothness of the marker motion (as in the technique that we exposed in 2.3.3.1), but their main focus is the validation of the technique. The third team [KL05] from GE Healthcare describes a method to detect and track the markerballs and register them with a linear transform. Their approach emphasises on the use of correlation with a marker template to detect and track them. The fourth team [ZHB⁺08] does not describe much the detection and tracking part but state that they use the classical linear transform for registration, a specific contrast enhancement based on the image histogram and a blur reduction filtering. Moreover, they restrict averaging only to the frames in a similar cardiac cycle phase, probably to enhance the sharpness of the DSE image. Finally, two other commercially available DSE products exist, one from Siemens called IC stent and one from Paieon called StentOp. Comparisons from an algorithmic standpoint were not possible as the technical details are not publicly available.

2.5.2.2 Comparison of the improvement in image quality

Two previous articles [Koo05, OSB⁺08] have evaluated the potential of DSE from an image quality improvement standpoint. The different teams had clinicians perform an assessment on a database of cine image sequences with the objective of comparing the images before and after enhancement. Formally, the individual clinicians were asked to rate each image according to a given scale, however the scales differ from one study to the other. Koolen [Koo05] used a 5-level scale defined by: (1) not visible, (2) poorly visible, (3) visible, (4) well visible, (5) crystal clear. Ohanessian [OSB⁺08] used a 4-level-scale defined as: (1) poor, (2) average, (3) good, (4) excellent. In the definition of these scales, the rater had the responsibility of choosing for each clinical case the details to focus the rating on. These scales share some similarities with the one that we have employed, but the main difference is in the precise definition of the levels. We made explicit the details to be visible at each level in hope to make the rating as reproducible and objective as possible. As expected, the different authors observed an improvement in image quality following the application of DSE. Koolen found that rates had increased from 1.7 ± 0.6 before to 3.5 ± 0.8 after enhancement. Ohanessian observed an increase from 2.03 ± 0.76 to $2.76 + / - 0.94$. The improvements are thus respectively of 1.7 and 0.9 points out of 5 (after compensation of Ohanessian results to emulate a 5-level scale). It is to be compared with the 1.7 point improvement we found. These positive results are in accordance with the noise reduction and contrast enhancement capabilities of DSE. It demonstrates that the potential inaccuracies of the registration that can impair the quality of the enhanced stent images are efficiently counterbalanced by the noise reduction and contrast enhancement.

2.5.2.3 On the clinical usefulness of DSE

One may wonder whether the demonstrated improvement in image quality end up in a clinical benefit. The study we led was not designed to support an objective assessment of this aspect since the enhanced images were not available during the procedures themselves. However, we relied on the experience of the readers to determine if the use of the enhanced images during the procedure would have increased the quality of the angioplasty. In our case, the two observers agreed that, in 13% of the cases, the enhanced image would have helped them to perform a better angioplasty. This question is very subjective and there is an important variability in the answers: the answer is positive in 40% of the cases for one observer and 19% for the second observer. Ohanesian [OSB⁺08] pursued a similar goal by asking to the observers about the usefulness of DSE compared with conventional cine images. They report a proportion of 72% positive answers. Despite the differences in the figures, both their study and ours tend to show that the increase in image quality due to the enhanced image is of clinical value in a number of cases. The exact percentage and typology of concerned cases is not known and is probably difficult to establish.

2.5.3 The role of DSE in clinical practice

As a consequence of the growing interest of clinicians for DSE, several studies have been conducted [MVP⁺07, Koo05, CSMS05, KJY⁺05, OSB⁺08, CAC⁺09, Cho10, FGBV09, MSB⁺10] that demonstrated that DSE is bringing a significant improvement in image quality and that it is deemed useful in clinical practice in a large proportion of cases ($\approx 70\%$). DSE can help evaluate the deployment of the stent and detect some under-deployments or non optimally treated lesions. In such cases, clinicians can post-dilate the stents to improve their expansion. The rate of stent post-dilation using DSE ranges from 30% to 70% depending on the study. Further studies are needed to better understand the role and the impact of DSE on clinical practice, for instance comparing post-dilatation rates with and without DSE. Only one study [Cho10] investigated the impact of DSE on the main clinical endpoints such as the occurrence of major adverse cardiac events and target lesion re-vascularisation. Based on a two arm set up involving nearly 900 patients, they demonstrated significant improvement of the clinical endpoints at mid-term follow-up. Moreover, various teams [Koo05, CSMS05, KJY⁺05, MVP⁺07, CAC⁺09] studied the relationship between the stent expansion quantification performed with this technique with respect to IVUS as gold standard. These five studies led to the interesting result that the relationship is linear with a rather good correlation ($r \approx 0.7$). It is a promising step towards quantitative assessment of stent expansion using DSE. However DSE alone cannot provide any information regarding the apposition of the stent onto

the vessel wall. The only recognized technique to perform this remains intra-vascular imaging.

2.5.4 Selected cases

We illustrate the interest of DSE on a set of clinical cases where it happens to be especially relevant. The first case we illustrate (see Fig 2.33) shows a deployed stent and a calcification. The harder plaque structure prevents the stent from optimally deploying and fully treating the lesion.




FIGURE 2.33: DSE image: Observe the plaque indicated by the arrow that prevents the proper deployment of the stent.

The second case we consider deals with stent positioning. When a stent is positioned in continuation of an existing one, good clinical practice consists in overlapping them at the junction with the minimal achievable overlapping area. This is extremely challenging, and stent visualization can facilitate this task. On the left image of Fig. 2.34 the stent undeployed is positioned without overlap with the existing one. The clinician rectified this before stent expansion in Fig. 2.34 right.

The third case deals with bifurcation stenting. A recent popular trend [MCTTY10] for these procedure consists in using a single stent in the main vessel that is inflated successively with different balloons to adjust its shape to the bifurcation and open the struts that jail the bifurcating vessel. In this context DSE can be used to visualize the modified shape of the stent. In Fig. 2.35 we can observe the uncommon shape of the stent resulting from this procedure.

Finally we present a rare but severe complication. In this case the patient presented at follow-up a coronary aneurysm at the very location where a stent had been deployed (Fig. 2.36). DSE at this location enabled to determine that the stent had fractured, damaging the artery. It allowed to understand the precise location and size of the orifice of the aneurysm. A specific treatment option (2 coils and 2 stents) has been performed to prevent the aneurysm from further inflating and possibly rupturing.




FIGURE 2.34: Use of DSE to optimize stent overlapping.

2.5.5 Hardware vs software DSE

The method discussed in this chapter relies on a standard Xray acquisition complemented by the use of specifically designed image processing algorithms. One may wonder if a specific Xray acquisition could produce the desired enhancement in stent CNR, without invoking image processing. To yield similar enhancement to the accumulation of 30 frames by the DSE algorithm, a single Xray shot shall deliver 30 times the conventional dose. It results on constraints on the tube and on the detector. The tube peak power shall support delivering such dose in less than 10ms (not to introduce blur), and the detector shall not saturate. It is approximately 6 times the capacity of an Xray tube, and 3 times the saturation level of a detector. None of these constraints can be fulfilled today nor in a near future.

2.5.6 Observed limitations of the technique

In 7% of the cases, stents do not appear clearly after stent enhancement (IQ scores equal to 1). There are various explanations that we can propose for this. The first one is that although the registration was perfect, the input CNR was so small that noise reduction was not efficient enough to make stent struts visible. Our set of test sequences includes on purpose approximately half of the sequences in challenging imaging conditions. Although this is much more challenging than daily clinical practice, we believe that it is not a good explanation for every case. Some of the low scores have been obtained in good imaging conditions. We do not have a final explanation, however we suspect that in a significant proportion of the cases there was registration error (Fig. 2.38 (a)). Since the landmarks are almost always very accurately detected by our method, we think that the registration errors are due to the fact that the stent does not follow the expected motion. We identified two cases. First, there can be motion between the stent and the landmarks since they are not rigidly linked. A visual inspection of the database




FIGURE 2.35: Use of DSE for a bifurcation treatment procedure. Top: bifurcation before treatment presenting significant stenoses. Center row: a first stent is deployed / a second stent is positioned next to it / both stent expanding (note the overlapping) / stent shape is further modified by balloon inflation in the side branch. Bottom: Final angiographic result of the treated bifurcation.

indicates that in some critical cases this motion can exceed 10 pixels. Our method does not behave well in these cases. Secondly, the motion of the stent occurs in 3D and may not be compensated by 2D image deformations. Fig. 2.37 shows a stent under two view points differing by a rotation of a few degrees around the long axis of the stent. We can observe that the appearance of the struts is completely different and therefore cannot be compensated by a 2D registration. A potential solution to these cases where stent registration is not perfect, could be to select a subset of images in the sequences in which the motion of the stent is less significant. Ross [RLM⁺05] noticed on one clinical case that using a subset of frames restricted to a specific interval of the heartbeat yielded a sharper stent image, similarly to Zaid [ZHB⁺08]. This research direction needs further studies to be confirmed.




FIGURE 2.36: Use of DSE to detect stent fracture. Left: a coronary aneurysm is indicated by the arrow. Center: DSE enables to see that the aneurysm is linked with the stent rupture indicated by the arrow. Right: after treatment.



FIGURE 2.37: Illustration of the variability of the appearance due to strut superimpositions patterns. (a) and (b) depict the same stent under two viewpoints differing from few degrees.

The registration model we used is also challenged in the case of multiple stents at a vessel bifurcation. In such a case there is typically one stent in the main vessel and another one in a side branch. The presence of the landmarks in one of the vessels is not enough to predict the motion of the other vessel. Therefore the second stent is usually blurred as a consequence of the incorrect registration. This clinical situation would require extending the method to detect and register landmarks in both vessels. Fig. 2.38 (b) shows an example of a result of stent enhancement at a bifurcation.

Another limitation to stent visualization is the presence of structures that are superimposed to the stent and hide it. The landmarks for instance, and especially the guide-wire, impair the visibility of some stent struts. Moreover, in some clinical situations, highly radio-opaque sternal wires that do not follow the motion of the stent can severely degrade its visualization. Specific algorithms [FLBV10] have been designed to address these limitations, the main one being discussed in the next chapter.




FIGURE 2.38: Illustration of imperfect registration in some cases (see text for details). (a) Although the landmarks are well registered and sharp, the stent struts do not appear and the upper part of the stent is blurry. (b) Case of two stents at a bifurcation. Only the marker-balls and the guide-wire supporting them are used as registration landmarks. Therefore, the stent in the side branch is not registered as accurately as the one in the main branch and is blurrier.

2.6 Conclusion

We have presented a comprehensive approach to DSE. We studied the problem from a clinical standpoint to define the requirements on the solution. The most original parts of our DSE technique, the guide-wire segmentation and non-linear registration, have been described and tested thoroughly. Validation on a large number of synthetic and clinical images demonstrated that DSE improves significantly image quality (by more than 1 point out of 5), works automatically (in 91% of the cases) and performs fast enough ($16.6 \pm 3s$) to be integrated in a typical angioplasty workflow with success. Moreover our experiments established that the use of the guide-wire as an input landmark to non-linear registration of images were beneficial to DSE in 28% of the cases. Further work may include addressing motion between the landmarks and the stent.

Chapter 3

Subtraction of the guide-wire in enhanced stent images

3.1 Abstract

The task of examining DSE images to determine whether the stent is correctly deployed requires precise visualization of the stent struts and assessment of the parallelism and straightness of the edges of the stent. However the presence of the guide-wire (though necessary for image registration) is not beneficial to this task (Fig. 3.1). In this chapter we detail a technique to remove the guide-wire from DSE images that enables the visualization of all the details of the stent even at the pixels locations where the guide-wire overlaps it. The technique is based on two properties of DSE images : their transparent nature (due the physics of Xray imaging) and the fact that the guide-wire absorption can be accurately approximated by a model whose parameters can be estimated directly from the images. We first introduce an innovative three-layer image model that has some advantages (the layers can be accurately estimated from one single image with low computational power), propose an algorithm that takes advantage of this layer model, and finally demonstrate the good performances of the method on clinical data.




FIGURE 3.1: (a) input image. (b) output image of StentViz "guide-wire subtracted", top regular DSE image, bottom guide-wire subtracted.

3.2 Introduction

Digital stent enhancement has been presented in Chapter 2. This image processing technique enables the visualization of the stent with conventional Xray in the cathlab. The presence of the guide-wire in the input image sequence has proven to be useful to accurately estimate the motion of the stent that is key to the sharpness of the DSE image. However, from a visualization perspective, the guide-wire has no interest and introduces difficulties. Indeed, the stent is a very fine mesh of wire of very low contrast whereas the guide-wire is of much larger contrast. According to Sections 1.2 and 2.4.1.4, the contrast of the guide-wire is approximately 5 times larger than the one of the stent. In the DSE image the guide-wire is extended through the stent and hides its struts whose visualization is relevant for the clinician. Moreover, when the guide-wire is partially stuck to the stent border, it hides it, potentially preventing the clinician from assessing its general shape. These two cases are illustrated on Fig. 3.2 and 3.3. Finally, it also limits our ability to set the contrast of the image on the fine details of the stent struts. Therefore we propose in this chapter to remove it from the images. We first present the model of transparent layer superposition and layer separation in Section 3.3. A method to perform guide-wire removal is exposed in Section 3.4. A gallery of results is presented in Section 3.5 and the benefits and limitations of the technique are discussed in Section 3.6. Finally conclusions are drawn.

Regarding image processing, our main contribution here is the use of a parametric model to solve a layer separation problem. From a clinical perspective, it is the first time a method is presented to remove guide-wire from DSE images. This technique is embedded into the StentViz software (General Electric) and used in routine clinical practice since 2010.

3.3 Background

In multimedia image processing, the task of removing the guide-wire would typically be addressed with inpainting techniques, marking the guide-wire pixels, and replacing their original values by computed ones according to some prior on the image [BSCB00, DCOY03, CPT03]. Typical priors include the regularity of the isophote lines [BSCB00] for instance. Although multimedia image editing techniques have been applied to X-ray images [BDN09], we note that there is a fundamental difference between both image types: in X-ray images the guide-wire does not occlude the stent. On the contrary, both objects superimpose. Therefore, on the pixels of the guide-wire, the stent information is also theoretically present (though much weaker than the guide-wire signal, thus almost

⁰The work presented in this chapter is covered by two patent applications [BV10, VB10] and was the subject of an oral presentation at the EuroPCR 2010 conference associated with an abstract published in EuroIntervention [FLBV10].



FIGURE 3.2: Scheme illustrating the difficulties in stent visualization introduced by the presence of the guide-wire. Top: the guide-wire hides the structure of the stent struts; Middle: the guide-wire hides partially the border of the stent; Bottom: ideal case for stent visualization where the guide-wire is not displayed (images are adapted from a Cypher stent photo [tim05]).




FIGURE 3.3: Stent visualization difficulties in clinical DSE images. Left: the guide-wire hides the structure of the stent struts; Right: the guide-wire hides partially the border of the stent.

invisible with a regular display). We aim at retrieving this information that is naturally present in the image rather than guessing a likely image content. This is the reason why we prefer to invoke the formalism of layer separation in transparent images rather than inpainting.

3.3.1 The transparent layer model for Xray images

Transparent images refer to images where the observed intensity of a pixel is the sum of the input intensities of the rays coming to the pixel attenuated by the integral of the attenuation along their paths according to the Beer-Lambert law [Bus02]. According to this model, if a ray r_j of input intensity Nj_0 , crosses a medium of attenuation $\mu(p)$ at a spatial position p , the output intensity $i(r_j)$ of the ray is:

$$i(r_j) = Nj_0 \cdot e^{-\int_{p \in r_j} \mu(p)} \quad (3.1)$$

An image formed with an Xray source is thus naturally transparent. Such images where one observes the transmitted light are often opposed to visible light photographs,

where one observes reflected light and where objects usually occlude one another. Transparent images are often modeled as superimpositions of layers [CAW02, ABL09, SI04, SI05, SABM04, TOM00, WA96, BBHP92, IRP94, SM91, ZLP⁺09]. In this simplified model, a layer is a plan on which each point is associated with an attenuation. The layers are usually represented without loss of generality as parallel planes (see Fig. 3.4). In this case, the image formation mechanism is modeled as X-rays crossing the layer plans. We denote I the Xray image, p a pixel location and $I(p)$ its intensity. According to this model, if a ray r_j of input intensity Nj_0 , crosses n planes at points of attenuation μ_i , the output intensity of the ray is according to the Beer Lambert law:

$$i(r_j(p)) = Nj_0 \cdot e^{-\sum_i \mu_i(p)} \quad (3.2)$$

The value $I(p)$ of the pixel a point p is the sum of the intensities of the rays reaching p :

$$I(p) = \sum_j i(r_j(p)) = \sum_j Nj_0 \cdot e^{-\sum_i \mu_i(p)} \quad (3.3)$$

Moreover, we generally assume that one ray only having a straight path from a point-wise X-ray source reaches a point-wise pixel detector. This model is simplistic and does not take into account many of the specificities of X-ray images such as scatter radiation or beam hardening for instance. However it is very practical to handle and thus commonly used in this context with satisfying results [CAW02, ABL09, SI04, SI05, SABM04, TOM00, WA96, BBHP92, IRP94, SM91, ZLP⁺09]. It can be seen as a basis for the Digital Subtracted Angiography (DSA) that is used daily in non cardiac imaging [MNV99]. In this model:

$$I(p) = N_0 \cdot e^{-\sum_i \mu_i(p)} \quad (3.4)$$




FIGURE 3.4: Illustration of the continuous (volumetric) Xray image formation model (left) and the discrete planar layer model (right).

More precisely, the transition from a 3D volume to a set of planar layers can be thought as a set of successive projections: Let us consider a 3D volume Ω through which the rays travel. We associate an attenuation factor $\mu(q)$ to every point q of this volume. The image intensity at a location p is the integral of the attenuation along the ray $r(p)$ reaching this pixel:

$$I(p) = N_0 \cdot e^{-\int_{q \in r(p)} \mu(q)} \quad (3.5)$$

Now let us consider a partition $\{\Omega_i\}_i$ of the volume Ω , the equation rewrites:

$$I(p) = N_0 \cdot e^{-\sum_i \int_{q \in r(p) \cap \Omega_i} \mu(q)} \quad (3.6)$$

We observe immediately that it is equivalent to an image formed with a layer model where the layers are defined by their attenuations $\{\mu_i\}_i$ following:

$$\mu_i(p) = \int_{q \in r(p) \cap \Omega_i} \mu(q) \quad (3.7)$$

This illustrates the generality of the layer model. A layer decomposition is thus any partition of the imaged volume. Generally the volume is split into layers that make sense from a physical/clinical point of view: organs or medical devices for instance. It is important to keep in mind that a planar layer is not the model of a planar object, but the model of a 3D volume as it appears under a given projection angle.

3.3.2 Layer separation in transparent images

We propose to consider the guide-wire subtraction task as a layer separation problem in X-ray images (which refers to decomposing an input image into the different theoretical layers composing it, *i.e.* estimating the layer images). This problem also arises in radiology for applications different from stent visualization where the layers correspond to different organs, bones, vessels or the interventional tools of a clinician (stents, balloons, catheters...). It is obviously an ill posed problem that consists in estimating many images (the layers) from a single one (the transparent image). Although it is an active research topic, current techniques only allow layer separation in some very special cases. Close [CAW02] proposed to process a sequence of moving images and to separate layers assuming that their motions are distinct (for instance to separate a static background layer from the moving layer of an artery of the heart). They proposed to estimate the motion of a layer, to produce a stabilized sequence where this layer is made static and to average the stabilized sequence. In the resulting image, the other layers that were not stabilized are blurred whereas the fixed layer is enhanced (denoised). Such a technique relies on strong hypotheses that limit its applicability:

- it assumes that the motions of the different layers are distinct. Therefore two layers undergoing the same motion cannot be separated.
- It assumes that layer motion can be estimated. This can be very challenging since the general intensity invariance hypothesis that is the basis of optical flow methods does not hold in transparent images. Therefore, motion estimation in transparent images is only feasible in a more reduced number of situations [ABL09, SABM04, TOM00].
- Moreover it assumes that the 3D motion of the layer can be compensated by 2D registration in the projection images.

Close *et al.* presented their general method with two main applications: vessel visualization enhancement and DSE. As previously mentioned in Chapter 2 this team was the first one to expose the concept of motion compensated averaging for DSE. DSE can therefore be seen as a layer separation process, where the layer regrouping the structures following the same motion as the landmarks is separated from the other ones. The strong hypotheses of the method of Close are generally well satisfied in this context, as previously discussed in Section 2.5:

- The stent is animated by the motion of the coronary segments it lies in, which is different from the background anatomy motion and the motion of the other coronary arteries.
- The motion of the stent layer can be estimated thanks to the presence of landmarks: guide-wire and balloon marker-balls.
- Motion compensation based on 2D non linear registration generally yields satisfying results.

However, since the stent and the landmarks have similar motion (it is the main hypothesis of chapter 2), they cannot be separated with this technique and as a consequence DSE produces an image where both are visible together.

All the other works on layer separation deal with applications that are different from DSE. The two most common ones are vessel layer separation in coronary angiograms and non medical multimedia applications. Sarel and Irani [SI04] proposed a technique to separate two layers when one of them is moving pseudo-periodically in a sequence of images and another method to separate layers when two images are available [SI05]. Several other methods have been proposed to separate layers when a sequence of images is processed. Shizawa and Mase [SM91] exposed a method in the 3D Fourier space for the special case where the layers move at a constant translational speed. Zhang et al [ZLP⁺09] separate the static background, the lungs and the coronary artery tree based

on scale and motion speed. Joint segmentation and motion estimation of transparent images has also been studied [WA96, BBHP92, IRP94, SI04].

Although it does not rely on the layer formalism, let us finally mention the large field of multi-energy acquisitions (dual energy being the most representative) that proposes solutions to separate objects in Xray images. It is based on the difference in their absorbance at different energies. It requires multiple acquisitions of the same static object with different Xray tube voltage [LAM⁺81]. This technique is thus only applicable to the limited number of images acquired under these conditions. For the application of guide-wire subtraction from enhanced stent images it suffers two major drawbacks: first, multi-energy would complicate the image acquisition procedure, second, the stent and the guide-wire are often composed of similar material (stainless steel).

Finally, DSA can be seen as a layer separation process. It consists in subtracting two images that only differ by one layer. It is commonly used to visualize vessels by subtracting a non injected and an injected image of the same anatomy. Therefore it requires two images.

To conclude this review of the literature, all known layer separation techniques require a set of several images. The case of the guide-wire layer that we are addressing has some particularities that make it solvable with one single image. We explain in the next section a dedicated technique to estimate it directly from the stent enhanced image.

3.4 Method

In the context of stent visualization we propose to consider a three layer model. The first layer is the anatomical background, the second the stent and the third the guide-wire. In such a model, we aim removing the guide-wire and background layers.




FIGURE 3.5: Planar layer modeling of a DSE image: 3 separate layers, namely the anatomical background, the guide-wire and the stent, are combined in the image formation mechanism.

3.4.1 Layer separation in DSE images

We first apply the log to the DSE image to have a linear relationship between the layers, where the value of a pixel is simply the sum of the values of the layers crossed by the Xray (Eq. 3.8).

$$I_1(p) = \ln(I(p)) = \ln(N_0) - \sum_i \mu_i(p) \quad (3.8)$$

$$I_1(p) = \ln(N_0) - (\mu_{background}(p) + \mu_{stent}(p) + \mu_{guide}(p)) \quad (3.9)$$

3.4.1.1 Background layer separation

Similarly to [ZLP⁺09] where a low resolution background is subtracted from the images, we separate the background based on scale characteristics. The background is typically composed of large structures that appear at a much lower scale than the objects of interest (guide-wire and stent). One can estimate the background by any state of the art technique, for instance low pass filtering [ZLP⁺09] or mathematical morphology approaches [Soi99, TZPB01, CAEB04, CM09, DKT05]. Once the background image has been computed, we subtract it from the original image. In our case this has already been performed in the DSE images thanks to the top-hat transform mentioned in Section 2.3.3.1. The resulting image I_2 only contains the stent and guide-wire information.

$$I_2(p) = \mu_{stent}(p) + \mu_{guide}(p) \quad (3.10)$$




FIGURE 3.6: Background subtraction in DSE images. Left: motion compensated integration of the images. Right same as left performed on background subtracted images.

3.4.1.2 Guide-wire layer separation

In order to constrain the ill-posed problem of estimating the guide-wire layer we introduce a-priori knowledge. We define a mathematical parametric model of the guide-wire and estimate it based on the DSE image. Injecting the estimated parameters inside the model allows the generation of the guide-wire layer, that is further subtracted from the

DSE image.

The guide-wire is accurately described by a curvilinear cylinder of circular section. Assuming that the angle formed between the incident X-rays and the guide-wire is constant over the piece of guide-wire that is considered, we can state that its absorpency is a function of the distance to its centerline only [BG09]. We want to emphasize that we do not assume that the guide-wire is parallel to the image plane. We only assume that it forms a constant angle with the imaging plane. It is a far less restrictive and more realistic hypothesis. The guide-wire section is typically composed of several concentric layers, the core (typically in stainless steel), coils or cover and coating composed of different materials [ENS⁺10] (see Fig. 3.7). Its tip, that is not the portion of interest in our case, has a different composition that make it more flexible and more radio-opaque.




FIGURE 3.7: Anatomy of a coronary guide-wire illustrating its main components: the core, the cover and the tip (Figure reproduced from [ENS⁺10]).

The model of the guide-wire thus requires the knowledge of two elements: the guide-wire centerline defined by a curve denoted C and the attenuation profile $\alpha(d_g)$ that maps a distance d_g from the guide-wire centerline to an attenuation (Eq. 3.11 and Fig. 3.8).

$$\mu_{guide}(p) = \alpha(d(p, C)) \quad (3.11)$$




FIGURE 3.8: Theoretical attenuation profile of a uniform cylindrical guide-wire imaged with a perfect system (MTF = dirac). Abscissa: distance to the centerline in pixels (the radius of the guide-wire is here 1,5 pixels). Ordinates: attenuation factor (no dimension).

In order to compute an estimate \hat{C} of C we segment the guide-wire in the image. Several techniques are suitable for this task. The one we use is detailed in Section 2.3.3.2. Once \hat{C} is known, we estimate α . We could argue that based on the exact composition of the guide-wire and of the imaging system response one could derive α . Although true from a theoretical standpoint, it turns out to be a difficult topic in practice. A precise estimation of α would have to be tuned to the particular composition of the layers composing each guide-wire, to the angle between the guide-wire and the incident rays (that is unknown), and to the precise system Mean Transfer Function (MTF) that mainly accounts for blurring and sampling in the physical imaging chain. This is reason why we rather propose to estimate α from the data available in the image. In order to estimate $\alpha(d_g)$ for a given d_g , we consider the set of values $\{I_2(p)\}_{p \in V(d_g)}$ of all the pixels in the image at distance d_g from \hat{C} with a tolerance δ (see Eq 3.12). Then we estimate a representative value from this set, for instance by computing the average or the median, that we assign to $\hat{\alpha}(d_g)$. In practice for d_g superior to a typical guide-wire radius, $\hat{\alpha}(d_g)$ is set to 0. Knowing \hat{C} and $\hat{\alpha}(d_g)$ we can compute the value of the guide wire at any pixel and thus compute the guide-wire layer:

$$V(d_g) = \{p \mid -\delta < d(p, C) - d_g < \delta\} \quad (3.12)$$

$$\hat{\alpha}(d_g) = \frac{1}{|V(d_g)|} \sum_{p \in V(d_g)} I_2(p) \text{ or } \text{Med}(I_2(p)) \text{ for } p \in V(d_g) \quad (3.13)$$

$$\hat{\mu}_{guide}(p) = \hat{\alpha}(d(p, \hat{C})) \quad (3.14)$$

The guide-wire layer is then subtracted from the image to produce the output image I' .

$$I' = \mu_{stent} + \mu_{guide} - \hat{\mu}_{guide} \quad (3.15)$$

In the resulting image the guide-wire is erased while preserving the image content. Fig. 3.9 and 3.10 show the application of the guide-wire subtraction between the two radio-opaque marker-balls and the estimated guide-wire layer.

3.4.2 Simultaneous display of two DSE images

3.4.2.1 General technique

The guide-wire subtraction procedure previously described generates an additional image over the classical DSE one. For safety reasons it is preferable to display both to the clinician, in case the guide-wire subtraction had created artifacts, or had removed some image content. However, from an ergonomical standpoint, having two images to review is cumbersome. The clinician would have to switch on the system from one to the other,




FIGURE 3.9: From left to right: DSE image, guide-wire subtracted image, estimated guide-wire layer, curve of the intensities of the pixels on the guide-wire in the DSE image. Observe on the curve all the variations induced by the stent signal as a consequence of the transparency of X-ray images.



FIGURE 3.10: From left to right: input image, estimated guide-wire layer, guide-wire subtracted image.

at the cost of extra time and interaction. Displaying both on screen, involves down-scanning the images by two to make them fit. It introduces a loss of resolution that is not acceptable (Fig. 3.11).




FIGURE 3.11: From left to right: regular DSE image, guide-wire subtracted DSE image, both downscanned to fit the display, both cropped on the area of interest to fit the display.

Such a problem could be tackled in multimedia images with sophisticated techniques like automatic image re-targeting [STR⁺05] that allows to resize an image while preserving the most important content. It is risky in medical imaging since it can induce false relationship/distances between anatomical features and potentially remove some relevant image content. Therefore we rather look for an optimal crop of the images, since it guarantees that the geometrical relationships are preserved. We take advantage of the fact that most of the time, when we want to display several images together we do not have to display them totally, but that only a given region in each image is of interest. We propose to process images automatically to determine this area of interest and to crop the images in a potentially non rectangular manner to make them fit into

one screen with no (or minimal) loss of resolution. For instance in an image of a coronary artery, the area of interest to the clinician is just a band of pixels around the artery, and the rest of the image can be cropped out. An additional difficulty appears then because the resulting images after cropping are not square nor rectangle anymore. However, in some particular cases, a set up can be found that allows to display them together with much less down-scan than if they had not been cropped. We illustrate this in the case of displaying two images of a same deployed stent on one screen.

Based on automatic processing, we infer the position of the stent, its privileged direction and its dimensions (length and diameter) that are similar in both images. In our case, for DSE images, we can rely on the marker-ball and guide-wire segmentation previously presented. It enables us to define in each of the two stent images an area of interest. Then we split the image for display into two areas (area A and area B) by a line passing through the center of the image and of same direction as the direction of the stents. In the first half of the image we will place the first stent so that none of its area of interest is out of area A. If necessary we will down-scan the first stent image to make its area of interest fit in area A. Similarly we will position the second stent image in area B. In order to maximize the zoom factor we position the image with the minimum down-scanning that allows the areas of interest to fit within the half images. From a practical standpoint, we decided to split the images only in a given set of four predetermined directions (horizontal, vertical, and the two diagonals) and to chose among these directions the one that is the closest to the direction of the stent. On Fig. 3.12 simultaneous display of stents is illustrated, choosing among the four predetermined orientations. The area of interest used for this display is illustrated on the image on the right, it is the union of the 3 colored squares.

3.4.2.2 Implementation details

The DSE images (with and without guide-wire subtraction) undergo up-scan and cropping before being displayed in one single image. There are two constraints to fulfill to produce the final image: the zoom factor shall be as close as possible to the one of the regular DSE image, and the stent shall be fully displayed. The area of interest is a band of pixels around the stent, that has been determined to be of approximately 8mm on each side of the guide-wire. It allows to display the full stent (whose diameter is typically 2.9 mm and at maximum 4.5 mm - see Table 1.1 and Section 2.4.2.1) and if necessary a part of bifurcating stent. Taking into account a typical conversion factor for cardiac interventions of 1.4 mm/pixel for the current design of the GE systems, this band shall be approximately 60 pixel wide on each side of the guide-wire (thus a total of $2 \times 60 = 120$ pixels). Our task consists in cropping the images while making sure that the area of interest is preserved. For the horizontal (resp. vertical) stents we can extract




FIGURE 3.12: Simultaneous display of two DSE image with/without guide-wire subtraction in the case of a diagonal stent. The area of interest is represented by the union of the three squares in the upper part.

the central horizontal (resp. vertical) band in the image. Retaining a band of 180 pixels enables matching the zoom factor of StentViz (2.8) and including the 120 pixel wide area of interest. For the diagonal cases it is a bit less straightforward. We test for zoom factors between 1.5 and 2.8 several different diagonal crops and retain the crop with the highest zoom factor where the 120 pixel wide area of interest is preserved. In practice we test with a simplified area of interest defined with a set of 3 squares, each 120 pixel wide, centered on both markers and the central point of the guide-wire (see Fig. 3.12).

3.5 Results

We display here a picture gallery of the results of the guide-wire subtraction in clinical conditions (Fig. 3.15 to 3.26). The overall level of the results looks good. Every figure is accompanied by a comment on the performance of the subtraction technique and when it is relevant on the clinical content. The method and many of these results have been presented at an oral presentation at the EuroPCR 2010 conference and are associated with an abstract published in EuroIntervention [FLBV10].

In order to compare the relevance of the median or the average estimators to estimate $\hat{\alpha}$ (Eq. 3.12) we also illustrate the guide-wire attenuation profiles estimated in 10 clinical sequences (Fig. 3.13). The estimated profiles with both estimators generally yield satisfactory results: The attenuation profile is globally decreasing with the distance to the guide-wire centerline, converging to null attenuation far enough from it. We note a minor numerical artifact: some curves slightly decrease beyond 0. This is due to the

fact that the level 0, defined versus the background that is itself estimated, contains a degree of uncertainty.

We observe that the highest curve seems to be different from the other ones. It can be explained by the corresponding image (Fig. 3.19) that contains an unusual amount of stent information, because of the particular procedure in which a new stent is deployed inside a pre-existing one. The additional amount stent metal tends to generate an over-estimation of the guide-wire attenuation in our method. We consider that the median performs better than the average on this case given that the curve converges to a value closer to 0. However the resulting images are visually almost indiscernible. This remark happens to be general over our dataset: in almost every case both estimators yield similar images. In one case only, the difference is noticeable (Fig. 3.14). The particularity of this case is that the guide-wire is at the same distance from the stent border in the whole segment of interest. Therefore the stent border pixels pollute the estimation at this very distance. The median happens to be more robust to this situation than the average. We thus recommend the use of the median estimator for task of estimating the guide-wire attenuation profile $\hat{\alpha}$.

Moreover we observe that the estimated guide-wire profiles are significantly different from the theoretical one presented in Fig. 3.8. We assume that it is mainly due to the system MTF: the profiles we estimate are likely to be low-pass versions of the theoretical ones.




FIGURE 3.13: Attenuation profiles of 10 guide-wires in clinical cases estimated with the average (left) and the median (right). Abscissa: distance to the guide-wire centerline in pixels. Ordinate: attenuation factor (no dimension). The average profile with both methods is plotted with a dashed line.

3.6 Discussion

3.6.1 The benefits of guide-wire subtraction

The first benefit of the guide-subtraction is the comfort it brings to assess stent deployment. The clear view of the stent, as illustrated in the previous image gallery, and the visualization of the stent struts and border without obstacle is helping the clinician. In




FIGURE 3.14: Guide-wire subtraction results: left average, right median. Observe on the image on the left the white line (artifact) at the center of the stent.

the cases where the guide-wire is stuck to the border of the stent, it can be of significant clinical interest. Such a case, in Fig. 3.27, depicts a non straight stent border beneath the guide-wire that was not straightforward to detect without the subtraction. In one of the cases that we have processed, illustrated on 3.18, the subtraction makes a large dark area appear behind a segment of the guide-wire that may be a calcified plaque. From a qualitative standpoint, we notice that the guide-wire subtraction brings more information in the cases where StentViz result is particularly good (IQ score 4 and 5 typically on the scale defined in Section 2.3.1). Indeed if the IQ is poor, there are less details to reveal in the narrow band of pixels concerned by the processing.

A potential benefit of guide-wire subtraction is easier and more optimal setting of the contrast of the image on the stent signal. Indeed without the guide-wire subtraction an image windowing that would be optimal for the stent would result in saturation all over the guide-wire. This benefit has not been much investigated in this study but is illustrated on Fig. 3.28. Finally it seems that the removal of the guide-wire that is a strong signal along a curve makes it easier to assess the parallelism of stent struts. This hypothesis would require further experiments to be demonstrated.

3.6.2 Position of the guide-wire subtraction in our processing chain

We have presented in this chapter a scheme where the guide-wire subtraction is the last piece of the DSE image processing chain. However the technique is general and applicable at other stages of the chain provided that the background layer has been removed. We assume that the results would be similar especially with the average estimator for $\hat{\alpha}$: The specificity of our non linear registration is to register the levelsets of the distance to the guide-wire among the images. Since it is exactly the sets that we compute in $V(d_g)$, the estimation of $\hat{\alpha}$ by averaging and the registration commute. Moreover it commutes also (obviously) with image combination by averaging. Therefore it would yield the same result to perform guide-wire subtraction right after background removal in each image independently or at the very end of our DSE processing. We prefer to perform it once only on the final DSE image to avoid unnecessary computations.

We note that it is not possible with a linear registration in the DSE algorithm. The guide-wires would not be registered and could therefore not be estimated in a similar manner. One would have to proceed to guide-wire subtraction before registration.

3.6.3 Limitations

We identified two root causes for the cases where guide-wire subtraction creates artifacts. The first one is that the guide-wire may not be accurately described by the parametric model we propose. More specifically, it can happen that its contrast is varying along its curve, whereas the model assumes that the intensity profile is the same at any point of the curve. This may be explained by variations of the angle of the guide-wire with the incident Xray beam as explained in [BG09]. The second cause of failure of the technique are errors in the estimation of the model from the data. In some cases the segmented guide-wire is not well centered on the actual one that results in artifacts after subtraction. In some other cases the estimation of the intensity profile results in an erroneous profile. Future work could include constraints on the intensity profile such as monotonicity, or the fit of a theoretical guide-wire profile model to the data observed in the image.

The layer separation technique presented here could moreover be extended to other components of the DSE image. In our approach we did not subtract the guide-wire layer outside of the segment delimited by the marker-balls because our estimation of the guide-wire centerline in this area. Therefore when the distance from the markers increases the subtraction creates artifacts (typically where the guide-wire is not accurately segmented/registered). The extension of the subtraction outside of the marker-ball segment would probably require an improved guide-wire segmentation in this area. Finally, parametric layer separation could also be extended in an attempt to subtract the marker-ball.

3.7 Conclusion

The framework of transparent layers has been applied successfully to the removal of guide-wire in enhanced stent images. It found here a clinically relevant application. The use of parametric models, that enables its resolution may be applied successfully to other situations. The technique described in this chapter is the basis of the second commercialized version of StentViz. The feedbacks from the clinicians about it are positive. They consider that in daily practice it facilitates the visualization of stents, and in some cases resolves ambiguities.




FIGURE 3.15: (a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.




FIGURE 3.16: (a) input image. (b) guide-wire subtracted image. The guide-wire subtraction is satisfactory but not perfect. We can observe in the upper section of the stent a residual guide-wire signal. Our hypothesis to explain it is that the guide-wire centerline estimation is less reliable in this area where the stent border and the guide-wire are very close one to another.




FIGURE 3.17: (a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction. The stent-signal is rather weak compared to the intensity of the guide-wire, therefore subtraction improves one's ability to assess the general shape of the stent.




FIGURE 3.18: (a) input image. (b) guide-wire subtracted image. The subtraction is not perfect in the lowest part of the stent where one can notice an unusually white signal. However, on the upper part of the stent, the subtraction drives the attention of the reader on a dark area. Its presence is also visible (but far less conspicuous) in the non-subtracted image. This area is likely to be a large calcification that can have impacted stent deployment in a plane perpendicular to the one of the image. Such a finding can encourage the cardiologist to repeat a StentViz acquisition with a different viewing angle.




FIGURE 3.19: (a) input image. (b) guide-wire subtracted image. The subtraction proved to be efficient in this stent in stent situation (two stents are superimposed in the lowest part). StentViz enables to understand the relative positioning of the two stents.




FIGURE 3.20: (a) input image. (b) guide-wire subtracted image. Perfect subtraction example.




FIGURE 3.21: (a) input image. (b) guide-wire subtracted image. Excellent guide-wire subtraction.




FIGURE 3.22: (a) input image. (b) guide-wire subtracted image. Excellent guide-wire subtraction.




FIGURE 3.23: (a) input image. (b) guide-wire subtracted image. Interesting case where the guide-wire is constantly at the same distance to the stent border. Thanks to the median estimator in the estimation of the attenuation profile the processing is robust to such situations.




FIGURE 3.24: (a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.




FIGURE 3.25: (a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.




FIGURE 3.26: (a) input image. (b) guide-wire subtracted image. Perfect guide-wire subtraction.




FIGURE 3.27: This case is of particular clinical interest. The guide-wire subtraction (on the right) reveals the border of the stent that is not straight, as indicated by the arrow. This finding was not obvious in the regular DSE image (left).




FIGURE 3.28: (From left to right: Regular contrast setting (no saturation); Aggressive contrast setting allows to improve the display of the stent but saturate the guide-wire; With guide-wire subtraction this contrast setting does not create saturation in the stent area.

Chapter 4

3D stent reconstruction

4.1 Abstract

Previous work has demonstrated the feasibility of 3D-stent reconstruction inside the cathlab through the processing of rotational acquisitions. In order to address the inconsistency in the projections due to the heart motion, this previous work relies on a motion compensated 3D reconstruction technique where detection of the balloon markers allows to infer the necessary image registration. In this chapter¹, we focus on several aspects of 3D stent reconstruction: automated marker detection and tracking, guide-wire subtraction, and registration. The challenge for the tracking algorithm is that the shape of the balloon is constantly changing within the image sequence due to the rotation of the camera. The key points of our method are the bottom-up approach, building primitives of increasing level of confidence, and the incorporation of the 3D properties of the object to track. The results on 22 clinical sequences are the best reported so far in this field. Moreover, the work we present regarding guide-wire subtraction and registration are new developments in the domain of 3D stent reconstruction.

4.2 Background

4.2.1 Clinical Background

The topic of stent visualization inside the cathlab has been extensively presented in Chapter 1. In Chapters 2 and 3 we presented technical developments to enhance the visibility of stents in projection images and subtract the guide-wire. These techniques, by improving the signal to noise ratio of stent images and removing disturbing structures, aim at producing the best possible 2D images out of a set of conventional Xray

¹Part of the work presented in this chapter has been done with Angela Chieh during her master's thesis internship. Quero agradecer a Angela por sua contribuição durante sua tese de Mestrado. I also want to thank Liliane Ramus for her collaboration and her support with the reconstruction and visualization softwares.




FIGURE 4.1: 3D coronary stent reconstruction at a bifurcation.

images. However, one of the key characteristics of stents, its three dimensional (3D) structure, has been neglected so far.

Conversely, intra-vascular imaging, like IVUS and OCT, presents the stent as a stack of 2D slices that better renders its 3D nature and enables the precise quantification of its expansion in every direction. For the sake of completeness, let us mention that Computed Tomography (CT) also enables 3D stent reconstruction. However, its temporal and spatial resolution does not allow to precisely visualize coronary stents.

In daily clinical practice, the 3D nature of the lesions and of the stents is systematically taken into account through repeated imaging under different viewpoints. Fig. 4.2 illustrates an under-deployed stent where only one of the two viewing angles allows to detect the defect. A potential solution to grasp the 3D nature of the stent from Xray images could be to repeat DSE under several viewpoints. However a 3D stent reconstruction is a more appealing objective and achieving such reconstruction is the aim of this chapter. Beyond the routine cases, a 3D stent reconstruction would be particularly useful in the complex ones, like bifurcations and overlapping stents, where the 3D stent shape is particularly complex.

4.2.2 State of the art

4.2.2.1 3D stent reconstruction

The first attempt to reconstruct coronary stents with conventional Xray has been performed by Close [CMW97], who demonstrated the feasibility of motion compensated stent reconstruction. The work that we present here is in the continuity of the work of




FIGURE 4.2: Illustration of a locally under-deployed stent in a pig artery. In image (a) the under-deployment circled in blue is visible. In image (b), under a different viewing angle the defect is hardly noticeable (images from [Per08]).

Perrenot [Per08] who developed during her PhD thesis at GE a 3D stent reconstruction technique exploiting the rotational acquisition mode available on the current imaging systems. This rotational acquisition called *cardiac spin* consists in a rotation of the C-arm in the axial plan with: CRA/CAU angle constant equal to 0° , and LAO/RAO angle spanning a range of 200° while the source to image distance is kept constant. The main difficulty addressed by Perrenot compared to conventional 3D reconstruction is that the stent is animated by the cardiac and respiratory motions in the set of 2D cone-beam projections. Similarly to DSE (Chapters 2 and 3), the markers of the balloon are used to infer the position of the stent. Then the reconstruction geometry is corrected to compensate the motion of the stent. The results she obtained are very encouraging. Fig. 4.3 presents a 3D stent reconstruction by Perrenot where we can clearly see the separate struts of the stent. The same stent reconstructed with a CT scanner is of far lower resolution (Fig. 4.3). From a theoretical standpoint Perrenot has demonstrated that the approximations made by correcting the motion of the stent only with 2D transforms are valid in the close vicinity of the markers and acceptable in the context of this clinical application. Two other teams investigated motion compensated 3D stent reconstruction, one from Philips Healthcare [SFL+08] and one from Siemens Healthcare [RLH11a]. The approach of the first team, very similar to the one of Perrenot, also compensates the marker motion based on affine transformation. However instead of encoding the motion in the projection matrices they register the 2D projections. The second team uses a different approach that computes a 3D+t marker model instead of an average 3D model.

3D stent reconstruction has undergone little validation so far. Perrenot [Per08] compared it with IVUS on a handful of cases from a qualitative perspective. Schoonenberg et al.

[SFL⁺09a] compared it quantitatively to IVUS on one case and to stent manufacturer dimension specification on 10 cases, demonstrating reasonable accuracy.




FIGURE 4.3: (a) A frame of the input sequence. (b) Visualization of the stent reconstruction with the method of Perrenot [Per08]. (c) Same stent reconstructed with a CT scan. Images extracted from [Per08].

4.2.2.2 Marker detection and tracking

The spin acquisition used to compute the 3D stent typically involves 150 frames. The manual detection of the balloon markers is clearly impossible in daily clinical practice and has to be automated. One of the goals of the present study is precisely to design and validate a fully automated marker detection algorithm. The challenges raised by this task are the following: first of all, the balloon appearance and its 2D apparent motion change significantly along the run, as the acquisition system is rotating. The balloon appearance commonly changes from two distinct markers to both markers superimposed. Moreover, it can be superimposed in some parts of the run to highly radio-opaque structures (catheters, or wires) hiding it. Finally a significant number of anatomical disturbing elements impairs its detection. These situations are illustrated in Fig. 4.4.

Schoonenberg et al [SFL⁺08] have exposed some results on automated marker detection and its impact on the quality of the reconstruction. They obtain 84% of correct detection but do not give much details about the algorithm they use. They report that this performance is enough to produce 3D reconstructions similar to the ones obtained by manual marker localisation.

4.2.2.3 Overview of the chapter

The chapter is structured as follows. We first present a fully automated marker tracking algorithm for rotational acquisitions (Sections 4.3.2). The specificities of the 2D DSE algorithms that we exposed in the previous chapters are extended to the 3D case in Sections 4.3.4 and 4.3.3. Results on a database of 22 cases are reported (Section 4.4)

and discussed (Section 4.6). Finally, conclusions are drawn.

Our main contributions are on the one hand the description of a fully automated algorithm for marker tracking in spin acquisition that yields, to the best of our knowledge, the best performance reported so far and, on the other hand, the extension to 3D of the specificities of our 2D DSE algorithm: non-linear registration and guide-wire subtraction.



FIGURE 4.4: Five images extracted from a rotational acquisition. Note the changes in balloon appearance and marker position. In the second image, one of the markers is overlapping the tip of the catheter, making it hardly visible. In the third image sternal wires appear, and in the fourth the markers overlap the guide-wire tip.

4.2.3 The necessity of a new algorithm

Given the similarities of the proposed problem with DSE, it is natural to evaluate the performance of the DSE detection and tracking scheme exposed in Chapter 2 on the spin datasets. We measure it in terms of True Positive (TP) - the percentage of frames where both markers have been correctly detected - False Positive (FP) - where a false pair is detected instead of the actual markers - and False Negative (FN) - where no pair is detected. There are no True Negative (TN) since the markers are present in every frame of the sequences. Therefore the relation $FN + TP + FP = 1$ holds. The results on 3 patients, are reported in Table 4.1. The algorithm fails in typically half of the frames which is far below the target that we set above 90% (similarly to DSE). This justifies the need for a specific algorithm. However, some bricks of the DSE algorithm that are relevant for this application are kept or adapted.

Patient	TP (%)	FP (%)
1	46	1
2	53	7
3	69	4

TABLE 4.1: Result of the detection and tracking algorithm designed for DSE.

4.3 Method

4.3.1 Clinical sequence database

The database is constituted of 22 clinical sequences acquired on 11 different patients. For each patient it contains a sequence with the balloon deflated and another one with the balloon inflated to reduce the motion between the balloon and the stent [Per08]. In the presentation of our data, we use the labels 1 to 11 for the sequences with the balloon inflated and 1a to 11a for their counterparts with the deflated balloon. The sequences have been acquired on an Innova 2000 system (GE Healthcare). The pixel size is 0.2 mm, the framerate 30 fps, and the angular coverage 200° . For evaluation purposes, we have manually marked, on each sequence, the position of each marker. It enables the computation of the TP, FP and FN rates.

4.3.2 Marker detection and tracking in rotational acquisitions

4.3.2.1 High-level architecture of the algorithm

The object that we want to detect and track is an angioplasty balloon delimited by two marker-balls at its extremities. Each marker is a small (5 pixel diameter) radio-opaque ellipsoid and both of them are linked by a guide-wire whose visibility varies significantly depending on its properties, on patient size and on the imaging angle. The 2D apparent properties of this balloon change dramatically along the image sequence as a consequence of the rotation of the image chain around the patient. The challenge is to build a robust and fully automated algorithm for the detection and tracking of the marker-balls. This algorithm shall output the precise position of each marker in every image of the sequence and shall label each marker to enable 3D reconstruction. It is very challenging to detect the correct markers among false alarms on each frame independently. Therefore we use the smoothness of the motion of the markers from frame to frame and the conservation of their appearance properties over some frames to achieve accurate detection among a set of pre-detected candidates. We design a bottom-up approach that is building primitives of increasing level of complexity and confidence: firstly points, then pairs of points, then short tracks of pairs and finally a full track of pairs. This single track obtained at the end of the processing is used to determine a 3D marker pair model and to compute the motion-compensated tomographic reconstruction of the stent. The first stages of marker detection are similar to the ones presented in Chapter 2 for DSE.

From a high level perspective, our detection and tracking scheme performs as follows: images undergo a pre-processing to ease the marker detection process, then a series of potential markers are detected in each frame, and are grouped into potential pairs of markers. Based on typical markerball motion parameters, a tracking algorithm regroups

the pairs into short tracks. Once the entire sequence has been processed the short tracks are merged in an iterative process that produces longer tracks until convergence to a single track representing most of the frames, or until no tracks can be harmoniously merged together. A refinement step is applied to the selected track in view of completing it and improving it. At this stage a 3D marker pair model is estimated and used to perform marker labeling and a final refinement that reduces FP and FN. The overall structure of the algorithm is illustrated in Fig. 4.5.




FIGURE 4.5: General overview of the marker tracking algorithm.

4.3.2.2 Pre-processing

Contrasted objects in the image include not only our object of interest but also a number of anatomical structures which create potential false alarms. The main properties of the marker-balls are their small size and their high radio-opacity (they are more radio-opaque than common anatomical structures). Background intensity variations as well as objects larger than the marker-balls are eliminated by a black top-hat transform (structuring element of size 9 like in Chapter 2). The details that remain after the transform are those that were, in the original image, darker than the background and smaller than the structuring element. Therefore, not only the marker balls are present in the top-hated image but also are the guide-wire, the stent, parts of the catheters, bone edges, sternal wires and noise. Among these features, the sternal wires are the main source of FP and degrade the performance of the algorithm. We designed a higher level process to remove them. They are easily identified based on their large radio-opacity and spatial extent.

4.3.2.3 Pair detection

In a first step we identify marker-ball candidates in the transformed image by selecting the relevant local minima. The relevance of a minimum is assessed based on its amplitude, on the local noise variance, and on the amplitude of other surrounding local minima. The marker candidates are then paired together based on geometric criteria (min and max distance between candidate points) and the resulting pairs are sorted according to the relevance of the markers composing them. Finally a predetermined number of pairs are retained in each frame. The thresholds are such that the ground truth markers are selected in more than 95% of the frames on clinical test sequences. In

the next stages, we explain how we use temporal coherence to choose within this set of selected pairs the one that actually contains the markers.

4.3.2.4 Short track construction

Tracking consists in grouping together marker pairs that appear consistently over frames. Because of the rotation during the acquisition, the appearance of the markers is very variable over the sequence, making tracking a more challenging task than in DSE. We propose at this stage to regroup marker pairs together over sets of frames where their appearance and motion properties are invariant enough. This first level tracking is a causal algorithm, that outputs a catalog of tracks which length is limited to some dozens of frames. The catalog is built by considering the succession of frames in chronological order². For each frame the algorithm updates the tracks based on the detected pairs. Each pair is associated to the track that is considered the most probable or is used to initiate a new track. The probability for a given pair to be the continuation of a track is estimated based on a set of cinematic parameters (marker segment variation in position and orientation, direction of the motion and speed), image based parameters (correlation of a neighborhood around the markers) and track related parameters (the number of appearances and absences of the considered track). The essential parameters of this probability function have been estimated on a large set of short clinical sequences. The tracking algorithm differs from the one of DSE to take into account specificities of rotational acquisitions: First of all, the distance between the markers can become zero when the viewing angle gets aligned with their axis during the rotation. Moreover the angle and length variations are far larger in the rotational acquisitions than in the DSE sequences. These two specificities have been taken into account by adapting the cost functions of the tracking. Finally rotational sequences are significantly longer - 30 frames for DSE, 150 for spin acquisitions. In order to have accurate short tracks at this stage, we forbid new detections to be associated with the tracks that have not been updated since a maximal number of frames.

4.3.2.5 Track merging

Presentation of the algorithm At this stage the correct markers typically appear in some of the short tracks stored in the tracking catalog among other tracks of false alarms. Statistics on the catalog show that in more than 50% of the cases the markers are scattered in several short tracks. In such cases it is necessary to merge tracks to regroup the markers. Therefore, we aim at merging together the tracks containing the

²We note that the tracking scheme that we selected is order dependent: its output can be different if the sequence is reversed.

markers to build a single long track. In order to differentiate marker tracks from false alarm tracks and drive the merge we rely on three properties:

- Marker tracks are generally longer than false alarm tracks,
- marker tracks are smoother than false alarms tracks (according to the parameters listed in the tracking),
- two consecutive marker tracks usually merge in a smoother way than two false alarms tracks or than markers and false alarms.

Algorithm The track-merging algorithm proceeds as follows: for each two tracks in the tracking catalog, it computes a score associated to the action of merging them. Then, it finds the two tracks that merge with the best score. If this score is higher than a threshold, it merges the two tracks together to form a new track. This new track is added to the track catalog and the two merged tracks are removed from it. Otherwise (if the score is too low) no relevant merge can be achieved within the catalog and the catalog is un-modified. An additional stopping criteria is added to this first one to stop it once a long track covers most the frames: the process of merging tracks is stopped as soon as a track covers more than 95%.

The track merging score computed for each pair of two tracks in the catalog is the combination of:

- a track confidence score that favors tracks that have high probability of being actual marker tracks,
- a continuity score that favors tracks that merge smoothly,
- a novelty score that helps fill the missing frames as much as possible.

The track confidence score is similar to the one used to determine the marker track among the catalog in DSE. It combines the number of frames of the track, the average absorption of the candidate markers, and the average correlation of the candidate markers. The novelty score is the percentage of missing frames that would be filled by performing the fusion. The track continuity score, a more technical matter, is defined hereafter.

Track continuity score Let us call $A = \{a_i, i \in I_A\}$ and $B = \{b_i, i \in I_B\}$ the two tracks considered for merging, where a_i and b_i are the pairs detected in the frames indexed in $I_{A,B} \subset \{1, 2, \dots, n\}$, and n is the number of frames of the sequence. We do not consider the merging of tracks if one is included in the other one ($I_A \subseteq I_B$ or $I_B \subseteq I_A$). We transform the problem of merging tracks into the one of merging non

overlapping tracks (see Fig. 4.6). To this end, let us consider the following reduced tracks: $A_r = \{a_i, i \in I_A \setminus (I_B \cap I_A)\}$ and $B_r = \{b_i, i \in I_B \setminus (I_B \cap I_A)\}$. We forced the merged track $M(A, B)$ to be either $M(A_r, B) = \{a_i, i \in I_{A_r}\} \cup \{b_i, i \in I_B\}$ or $M(B_r, A) = \{a_i, i \in I_A\} \cup \{b_i, i \in I_{B_r}\}$ (*i.e.* the merge of two non overlapping tracks).




FIGURE 4.6: Merging two potentially overlapping tracks A and B is reduced to merging non overlapping ones by assigning all the common indexes to either one or the other.

We then define the set of couples of points $\xi(C, D)$, where two non-overlapping tracks C and D merge:

$$\xi(C, D) = \{(c_i, d_j) | c_i \in C, d_j \in D \text{ and }]\min(i, j), \max(i, j)[\cap (I_C \cup I_D) = \emptyset\} \quad (4.1)$$




FIGURE 4.7: Two non overlapping tracks C and D are represented. The merging point couples ($\xi(C, D)$ in the text) are represented with a wider black border.

Fig. 4.7 illustrates these points. For each merging point $(c_i, d_j) \in \xi(C, D)$ we define a local continuity score $\omega(c_i, d_j)$ that is a function of the following parameters:

- the gap between the two tracks: $|i - j|$.
- the distance d between the centers of the pairs.
- the variation of length of the pairs Δl .
- the variation of angle $\Delta \theta$.
- the correlation c between patches centered on the extremities of the pairs.

We compute it with the following formula, where the σ_x are sensitivity parameters that have been set on our test database:

$$\omega(c_i, d_j) = 1 - e^{-\left(\frac{|i-j|}{\sigma_g} + \frac{d}{\sigma_d} + \frac{\Delta t}{\sigma_t} + \frac{\cos(\Delta\theta) - 1}{\sigma_\theta} + \frac{c}{\sigma_c}\right)} \quad (4.2)$$

An overall continuity score is computed as:

$$\bar{\omega}(C, D) = \frac{1}{\#\xi(C, D)} \sum_{(c_i, d_j) \in \xi(C, D)} \omega(c_i, d_j) \quad (4.3)$$

4.3.2.6 Determination of the marker track and refinement of the track

At this stage the tracking catalog contains a set of long tracks. In order to identify the one that is the most likely to contain the actual markers we use a metric similar to the confidence metric presented in the last section. The track that raises the best score is selected. The confidence in the track selection is high and the selected track is the only one that will be considered for the remainder of the algorithm. It typically contains most of the markers ($\approx 85\%$) of the sequence but lacks some markers and includes a small proportion (typically a few percents) of false alarms. This is below the target we aim at, (see Section 4.2.3) and probably suboptimal for 3D reconstruction. The refinement step that we describe hereafter aims at solving these issues. The first part of the refinement identifies the frames where wrong markers (false positives) were detected. The track being known, its typical motion parameters (namely frame-to-frame variation of length, orientation and position) can be precisely estimated. The frames where the variation are statistically unlikely are removed. The second part of the refinement detects the markers in the frames where they were not previously detected. The fundamental hypothesis guaranteed by the acquisition protocol is that the markers are present in every frame. In a given frame where the markers were not found, the search area and the pairing criteria are refined based on the known position of the markers in the past and future frames. It allows detecting marker pairs that may have been lost in the first stages of the algorithm or where FP have been rejected. The refinement procedure is iterated until convergence or until a maximum number of iterations is reached. FP and FN reduction techniques are presented thoroughly hereafter.

False positive reduction Fig. 4.8 illustrates the evolution of the length between the candidate markers in a track. We can observe peaks in the interval of frames [150, 155] and [58, 160] that correspond to false positives in the frames 152, 153 and 159.

This suggests that abnormal variations in the evolution of the parameters describing the markers can be used to detect false positives. Therefore, for every description parameter p (length, angle and position) we define a measure at each frame $\delta(p)[i]$. We compute a smoothed version of this measure by local averaging on a window of a few

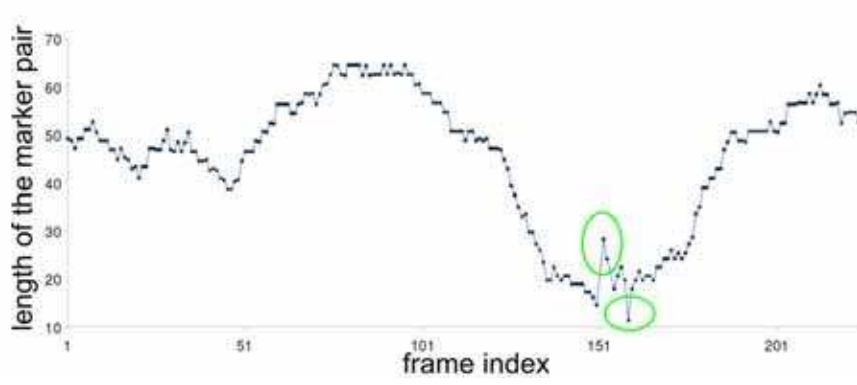


FIGURE 4.8: Plot of the length of the detected marker segment versus the frame index. False positives are circled in green.

frames and the difference between the original measure and its smoothed version, that we denote $D(p)[i]$. It enables us to compute local measure variations (which is basically the high frequency part of the measure). To identify unlikely peaks in the measure variation, we estimate its average $\mu(D(p))$ and standard deviation $\sigma(D(p))$ over the whole set of frames. Given a sensibility parameter k_p , we define a normalized measurement $\hat{D}(p)[i]$:

$$\hat{D}(p)[i] = \frac{D(p)[i] - \mu(D(p))}{k_p \sigma(D(p))} \quad (4.4)$$

We gather the normalized measurements for all the parameters length l , angle θ , position x and y in a vector $\hat{\mathbf{D}}[i]$:

$$\hat{\mathbf{D}}[i] = \begin{pmatrix} \hat{D}_l[i] \\ \hat{D}_\theta[i] \\ \hat{D}_x[i] \\ \hat{D}_y[i] \end{pmatrix}$$

We reject a frame (considered false positive) if:

$$\|\hat{\mathbf{D}}[i]\|_\infty \geq 1. \quad (4.5)$$

The philosophy of this test is to remove detections that are unlikely to belong to the distribution of measurements. This test is similar in its form (threshold of the type $\mu + k\sigma$) to typical statistical tests. We observed that the method gives satisfying results by using the second derivative as a measure of the parameter. It can be explained by the fact that the parameter itself and its first derivative are put under control by the tracking. The sensitivity parameters k_p are typically of the order of 5.

False negative reduction The chosen track originally contains some FN. Their proportion is most of the time increased by the FP reduction that transforms the detected FP into FN. This motivates the design of an FN reduction block following the FP reduction one. We refine the pair extraction for the FP frames by considering an ROI adjusted based on the detections in the closest previous and next frames (resp. indexed i_p and i_f). In this ROI we set less restrictive pairing constraints to increase the likelihood that the actual marker pair is selected. The main hypothesis in the definition of the ROI, is that the motion between frames i_p and i_f is close to a monotonous motion in both x and y coordinates. It is a good approximation given the trajectory of the imaging system. The main sources of deviation with respect to this assumption are the cardiac and respiratory motions. Therefore the ROI for frame i is defined by:

$$ROI(i) = \{(x, y) \mid x_{min} \leq x \leq x_{max} \text{ and } y_{min} \leq y \leq y_{max}\} \quad (4.6)$$

where

$$\begin{aligned} x_{min} &= \min(x(M_{1,i_p}), x(M_{2,i_p}), x(M_{1,i_f}), x(M_{2,i_f})) - B(i, i_p, i_f), \\ y_{min} &= \min(y(M_{1,i_p}), y(M_{2,i_p}), y(M_{1,i_f}), y(M_{2,i_f})) - B(i, i_p, i_f), \\ x_{max} &= \max(x(M_{1,i_p}), x(M_{2,i_p}), x(M_{1,i_f}), x(M_{2,i_f})) + B(i, i_g, i_d), \\ y_{max} &= \max(y(M_{1,i_p}), y(M_{2,i_p}), y(M_{1,i_f}), y(M_{2,i_f})) + B(i, i_g, i_d), \end{aligned} \quad (4.7)$$

and $B(\cdot)$ is an additional margin that enables handling motion further than the bounds dictated by a strictly monotonous model. It is defined by:

$$B(i, i_g, i_d) = b + c \cdot \min(i - i_g - 1, i_d - i - 1), \quad (4.8)$$

where b and c are constants that have been set based on our database.

Fig. 4.9 illustrates this process. We can observe the ROI and the list of pairs. Since the region where the processing occurs is reduced owing to the knowledge of the position of the markers in previous and next frames, and the parameters of the pairing relaxed, it is highly likely that actual markers are among the pairs. In this example, that actual marker pair was not present in the original pair list. We can observe that it is present in the new one (Fig. 4.9(e)).

The best pair is chosen according to a continuity score is similar to the one defined in Section 4.3.2.5. As in the tracking process the best pair is added to the track only if its score is greater than a minimal threshold.




FIGURE 4.9: (a) et (b): Frames 158 and 160 of a clinical sequence. The actual markers have been detected at the locations indicated in green. However, in the intermediate frame (number 159) they have not been detected. (c) The set of detected pairs in frame 159. The actual marker pair does not belong to them. (d) The ROI for pairing and the new set of pairs in frame 159, based on the detected markers in the previous and next frames. (e) Extract from (c) where we can observe that the actual markers did not belong to the set of detected pairs.

4.3.2.7 Track refinement

The process of reducing false positives followed by false negatives is beneficial to the quality of the track. Therefore iterating it generally further improves the track. In practice we iterate this process until the track stays unchanged or a maximal number of iterations is reached.

4.3.2.8 Labeling of the markers

The last challenge of the marker tracking is the labeling of the markers *i.e.* identifying all along the sequence the proximal and the distal markers. This step is essential to avoid performing an inverted registration during the motion compensated tomographic reconstruction. Approaches based on the analysis of their relative position along the sequences fail because the two markers may superimpose and invert their relative positions (see Fig. 4.10). As a consequence, we have adopted a 3D geometrically oriented approach. First of all, the center and the direction of the segment forming the markers in 3D are roughly estimated in a mean least square fashion. This rough estimate corresponds to an average position in the sequence. It is then projected on each frame and the position of the markers versus the segment and the center allows their labeling. We define these operations specifically hereafter.

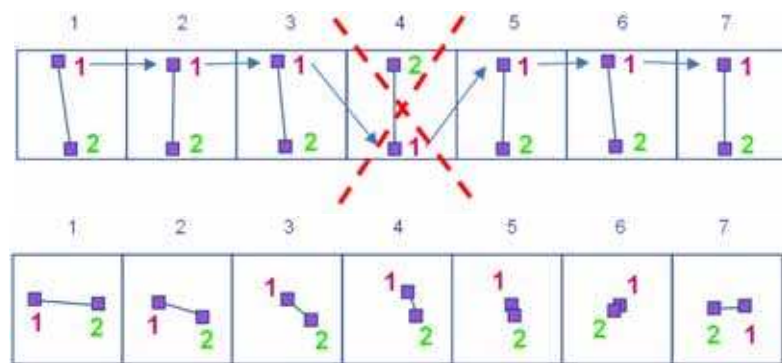


FIGURE 4.10: Top: Example of wrong marker labeling. Bottom: A typical case where marker labeling is a difficult case because the markers invert their relative positions along the sequence.

4.3.2.9 3D reconstruction of the marker line segment

We first reconstruct the average 3D position of the center point of the markers. The 2D position of the center point is computed thanks to the position of the markers in the track. We consider the set E_d of the back-projection lines of this point over all the frames. We

define the average position, \hat{p} as:

$$\hat{p} = \operatorname{argmin}_{p \in R^3} \sum_{D \in E_d} d(p, D)^2 \quad (4.9)$$

where d is the Euclidean distance function. It is illustrated on Fig. 4.11.

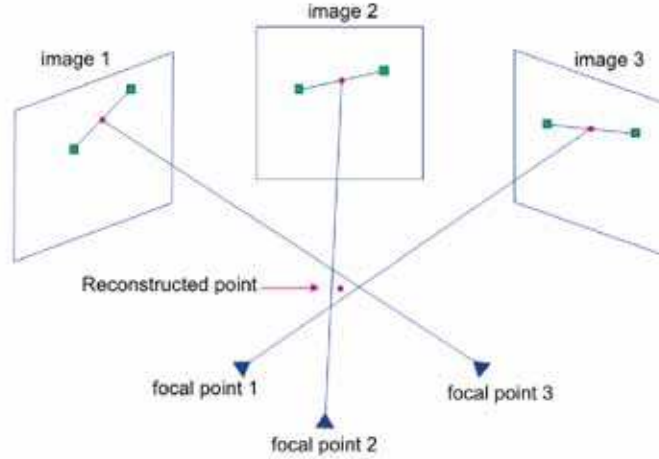


FIGURE 4.11: 3D reconstruction of the center point from 3 detections.

Once the center point has been reconstructed, we determine the direction of the line segment joining the markers. To this end, we consider the set of planes formed by the triplet of points: $(F_i, M_{1,i}, M_{2,i})$ in each frame i , where F_i is the focal point, and the other two the markers in the image plane. The normal vectors $\{\mathbf{n}_i\}_{i < n}$ of these n planes, illustrated in Fig. 4.12, are computed according to:

$$\mathbf{n}_i = \frac{\overrightarrow{F_i M_{1,i}} \wedge \overrightarrow{F_i M_{2,i}}}{\|\overrightarrow{F_i M_{1,i}} \wedge \overrightarrow{F_i M_{2,i}}\|}, \quad (4.10)$$




FIGURE 4.12: Plane associated with a marker couple i .

We express the computation of the direction as a mean square minimization problem, searching for a unitary vector \mathbf{d} such that:

$$\mathbf{d} = \underset{\mathbf{v} \in R^3 \text{ and } \|\mathbf{v}\|=1}{\operatorname{argmin}} \sum_{i=1}^N (\mathbf{v} \cdot \mathbf{n}_i)^2 \quad (4.11)$$

Let us consider the following matrix A :

$$A = \begin{pmatrix} n_{x_1} & n_{y_1} & n_{z_1} \\ & \vdots & \\ n_{x_N} & n_{y_N} & n_{z_N} \end{pmatrix}$$

It is well known that the eigenvector d of $A^\top A$ associated with its minimal eigenvalue is a solution to Eq 4.11. A rough 3D reconstruction $(\tilde{M}_{1,3D}, \tilde{M}_{2,3D})$ of the marker segment is then computed by positioning the two markers at an arbitrary distance from the center point along the 3D segment direction:

$$\tilde{M}_{1,3D} = \hat{p} + k \mathbf{d} \quad (4.12)$$

$$\tilde{M}_{2,3D} = \hat{p} - k \mathbf{d} \quad (4.13)$$

4.3.2.10 Application to labeling and track refinement

The labeling of the markers in the projections is performed thanks to the rough 3D model $(\tilde{M}_{1,3D}, \tilde{M}_{2,3D})$. Both markers are projected on every frame and their center is aligned with the center of the detected marker pair. Then, label i is associated with the closest marker to $\tilde{M}_{i,3D}$, as illustrated on Fig. 4.13.

We further refine the track to minimize FN and FP iteratively. At this stage the FN reduction can benefit from the 3D information. Given the labeling of the markers, we reconstruct them in 3D using the same technique as for the center point. It yields a 3D estimate of the average position of the markers $(\bar{M}_{1,3D}, \bar{M}_{2,3D})$. Then, during FN reduction, each pair is evaluated versus the projected 3D model in terms of angle and length to decide upon the one to select. This ends the marker tracking.

4.3.3 Guide-wire subtraction

Similarly to the 2D DSE case (see Chapter 3), the guide-wire that is reconstructed has no clinical added value and induces visualization difficulties. In the context of volume rendering the guide-wire occludes the stent struts that are behind it in the scene. Moreover, when it gets close to the stent border it prevents the user from properly visualizing the general shape of the stent. Therefore, we propose to remove it. There are basically two ways to deal with this task: processing the 3D volume or pre-processing the projections before reconstruction. According the work exposed in Chapter 3, we favor the second approach. We detect the guide-wire between the markers in each frame




FIGURE 4.13: Projection of the rough 3D model on the image, translation of the projected segment and labeling of the detected marker couple (in green).

of the input spin image sequence thanks to the technique presented in Section 2.3.3.1. This algorithm demonstrated its robustness in this situation, achieving a successful segmentation in more than 95% of the images. Then guide-wire subtraction is performed according to the parametric model developed in Chapter 3. Given the properties of the algorithm that is based on the transparent nature of Xray images, the fundamental hypotheses of tomography are not violated by the processing. After subtraction, the observed absorbance at the location where the guide-wire was is theoretically exactly the one that would be observed if it were not present. Fig. 4.14 illustrates guide-wire segmentation and subtraction in rotational acquisitions.

4.3.4 On the choice of the registration for 3D stent reconstruction

The method proposed by Perrenot [Per08] for 3D stent reconstruction is based on the motion of the markers. Once the markers have been detected in every image of the input sequence an average 3D marker pair model is reconstructed. This model is forward projected on the images of the spin and the difference between the detected and projected positions of the markers is compensated by modifying the projection matrices. Then standard cone beam 3D reconstruction is performed. Similarly Schoonenberg *et. al.* [SF09] detects the position of the markers in every image of the spin acquisition and reconstructs a 3D marker pair model that is projected on each frame of the spin. However instead of encoding the position differences in the projection matrices, the authors warp the projection images so that the detected markers match the projected model before




FIGURE 4.14: 3 images of a spin sequence. Top row: original image. Central row: segmented guide-wire. Bottom row: subtracted guide-wire.

performing a 3D reconstruction. In both cases the motion compensation is performed with a similarity transform. The second approach can be easily extended to different registration techniques. It can be formalized through the following scheme:

1. Detect landmarks in the 2D projections.
2. Compute 3D landmark model.
3. Project the 3D landmark onto the 2D projections.
4. Register the 2D projections so that the detected landmarks match their projected 3D model.
5. Perform cone beam 3D reconstruction.

We note that the approach of Perrenot [Per08] does not necessitate to explicitly transform the images and thus avoids the interpolation associated with this kind of techniques. It is an advantage since it allows preserving the sharpness of the original images.

4.3.4.1 Linear constant radius registration

We propose here to modify the 4th step of the Schoonenberg *et. al.* scheme. The classical marker based registration (a similarity transform) suffers drawbacks that we have discussed in Chapter 2. We proposed here an alternative one that combines the smoothness of the linear transformation with approximate stent radius preservation. The main idea behind its design, is to account for the change in distance between the markers by performing scaling only in the direction of the marker segment, thus not modifying the diameter of the stent (see Fig. 4.16).

The registration step aims at compensating the differences between the projected and detected markers. From a geometrical standpoint, the differences between the length of the detected and of the projected marker segments can be attributed to foreshortening or to global image scaling. Foreshortening is due to a change in orientation versus the Xray beam, whereas global scaling is due to a change in the source to object distance. These changes are due to the respiratory and cardiac motions, since the system geometry is taken into account while projecting the 3D marker model. The change in source to object distance, however is supposed to have little impact on the marker segment length. A typical displacement of 1cm results only in a 1% variation of the marker segment. For a typical stent length of 20mm that translates to 140 pixels the change is only of 1,4 pixel (below the radius of the markerball). Therefore we believe that when a significant change in marker segment is observed it is to be attributed to foreshortening rather than global scaling. Under this hypothesis, the correct way to account for a difference between the length of the projected and detected marker segment is to stretch the image only along the direction of this segment followed by a rigid transformation. We illustrate this in Fig. 4.15. In this example the stent has an out of plane rotational motion around (Ox) from the position of (a) to (b). If the change in distance between the markers is compensated by a global scaling, its factor $\lambda = \frac{d_1}{d_2}$ also affects the diameter of the stent, resulting in a false diameter change (Fig. 4.15 (c)). With the proposed technique the stent diameter is nearly unchanged. Its preservation is crucial to the quality of the reconstructed 3D stent.

The scaling in the direction of the markers is modifying the image I to produce an image $\text{Stretch}(I)$ according to:

$$\text{Stretch}(I) = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} \lambda & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} I \quad (4.14)$$

Where θ is the angle formed by the marker segment and the abscissa axis. The rigid transform T that maps the markers of $\text{Stretch}(I)$ to the projected ones is:

$$T(X) = \begin{bmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{bmatrix} X + \begin{bmatrix} X_0 \\ Y_0 \end{bmatrix} \quad (4.15)$$

where α is the angle between the projected and the detected marker segments and (X_0, Y_0) the translation between their centers. The overall registration yields a registered image I_r defined by:

$$I_r = T(\text{Stretch}(I)) \quad (4.16)$$

We note that $T(\cdot)$ is a rigid transform (preservation of distances) since the stretching has already been compensated by $\text{Stretch}(\cdot)$. Moreover, the overall transform coincides with the guide-wire based transform (Section 2.3.3.4) when the guide-wire is a straight line (this is not the case of the classical affine transform).




FIGURE 4.15: A situation where affine registration is not optimal. Between images (a) and (b) we assume that the marker segment has undergone a rotation out of the image plane, around the (Ox) axis. We illustrate in red the marker segment of image (a). Note the change in the marker segment length between figures (a) and (b). The change in distance between the markers in images (a) and (b) is compensated by a global scaling with the similarity transform that modifies the (nearly unchanged) diameter of the stent (image (c)). In image (d) a stretching along the axis of the markers but not perpendicularly keeps unchanged the stent diameter while registering the markers. In these two last images the registered markers are illustrated in green.

Finally, we note that this linear registration can also be embedded in the framework of Perrenot [Per08] that directly modifies the projection matrices. We would thus avoid the sharpness loss resulting from the image resampling during registration. In the results section we present examples with this transformation and with the similarity transform to register the detected marker on the projected marker model.




FIGURE 4.16: Motion compensated images. Left and center: Two registered images in a phase where the gantry is static. Right: Registration of the manually segmented stents. Top row state of the art similarity transform, bottom row linear diameter preservation registration. Observe that the registered stents match better with the proposed registration than with the classical affine one (as indicated by the arrows).

4.3.4.2 Guide-wire based registration

We consider here a non-linear warping accounting for the guide-wire deformation similarly to what we have exposed in the 2D case in Section 2.3.3.4. The 3D motion compensated scheme is thus modified by extending the concept of landmarks to the markers and the guide-wire and by modifying the registration process. It can be rewritten as follows:

- Detect the markers in the 2D projections.
- Detect the guide-wire in the 2D projections.
- Compute an average 3D landmark model (including the markers and the guide-wire).
- Project the landmark model onto the 2D projections.
- Non-linearly register the detected landmark on the projected landmark model.
- Perform cone beam 3D reconstruction.

In order to compute a single point of the guide-wire model, one has to determine its position in every projection image. However, this is not completely straightforward,

since we do not know the correspondences between the points of the guide-wire in the different images (at the exception of the markers that are clearly identified). A solution would be to use the epi-polar lines to match points and account for frame to frame motion (see [BMVA06] for instance). We propose an alternate simple yet efficient approximation using the arc-length along the guide-wire. In order to account for the arc-length variations due to the projection angle in the spin acquisition we normalize it by the arc-length between the two markers. We then assume that points of normalized arc-length correspond to the same physical guide-wire point along the sequence. It is equivalent to making the hypothesis that the angle between the guide-wire and the Xray beam is constant between the markers. This approximation can be grounded by the small length of the guide-wire segment (making conic projection similar to parallel) and neglecting the guide-wire curvature. Therefore we sampled each segmented guide-wire with a constant step in normalized arc-length to generate a set of corresponding points. Each point is then reconstructed in a least square fashion similarly to Eq. 4.9.

Once the average guide-wire has been reconstructed in 3D, it is projected back onto the images. However, this guide-wire is typically a polygonal structure, since only a set of sampled points have been reconstructed. Since the registration technique exposed in Section 2.3.3.4 assumes that the curve is differentiable, a fit is performed on the projected guide-wire points to yield a smooth parametric structure. Finally this structure is used as input to the non-linear registration.

4.4 Results

4.4.1 Marker detection and tracking

The algorithm described above relies on several figures of merits in the different steps (selecting pairs, building tracks, merging tracks ...). To address the problems of selecting the numerical parameters involved in these different figures of merits, we did our best to identify linked parameters and limit at maximum the parameter sets. We then collected a database of sequences either static (*i.e.* DSE sequences without rotating the image chain around the patient) or dynamic (rotational acquisition). Manual identification of the markers has been performed in these sequences. The parameters mentioned above have been adjusted by optimizing the results on this database. In this framework, “results” means essentially percentage of correctly identified markers. At the end, the algorithm was applied on 22 clinical sequences of average length 145 frames (for a total of nearly 3200 frames). In the table below, we report the percentage of frames where the actual markers are detected and the percentage of frames where a false alarm is detected at different stages of the algorithm. We can observe how each brick improves




FIGURE 4.17: Illustration of the non-linear registration and of the guide-wire subtraction on three images taken at different instants during a rotational acquisition. We represent the segmented guide-wire in green, and the projected 3D guide-wire model in red. Top row: original images. Bottom row: registered and guide-wire subtracted images. Left: We can observe the discontinuity of the guide-wire based registration far from the segmented guide-segment (notice the “broken guide-wire effect”). Middle: the projected and detected guide-wires occasionally match very closely, typically once per heart beat, as in this example. Right: notice the difference in curvature between the projected and detected guide-wires.

the quality of the track. The results listed in table 4.2 demonstrate that the level of performance of our approach is satisfactory. In 90% of the sequences (20/22) more than 97% of the markers are detected. In the worst case 94% of the markers are detected. In average the TP rate is $98.5\% \pm 1.6$, the FP $0.6\% \pm 1.1$ and the FN $0.9\% \pm 1.0$. Finally there is absolutely no labeling error. These results are very promising. The detected markers were used to make a 3D motion compensated reconstruction of the stents. The output 3D volumes were qualitatively compared to the volumes obtained with a manual segmentation of the markers. No visual difference was noticed. From an applicative standpoint these results can be considered as excellent.

4.4.2 Detailed results of the different algorithms

The algorithm that best improves the TP rate is the track merging. It raises it on average from 82% to 95%. Moreover the improvement of the minimum performance is more striking than the average one: It improves it from 55% to 85% of TP, *i.e.* from results that could not yield a satisfactory 3D reconstruction to results enabling it. However, the FP is increased in 18% (4/22) of the cases in our database, justifying the need for a specific FP reduction brick. Merging has been performed in 59% of the cases (13/22),

Patient	<i>Tracking</i>		Merge		FP and FN reduction		Final		
	TP	FP	TP	FP	TP	FP	TP	FP	FN
1	61,6	3,8	91,2	3,8	97,5	0,0	100,0	0,0	0
2	55,1	1,3	96,2	1,3	96,8	1,9	97,5	1,9	0,6
3	76,4	0,6	96,9	0,6	98,8	0,0	98,8	0,0	1,2
4	88,7	8,5	91,5	8,5	93,0	4,9	94,4	3,5	2,1
5	97,3	1,3	97,3	1,3	97,3	0,0	98,7	0,0	1,3
6	98,1	0,0	98,1	0,0	100,0	0,0	100,0	0,0	0
7	99,4	0,0	99,4	0,0	100,0	0,0	100,0	0,0	0
8	55,6	9,4	85,0	11,3	88,1	5,0	99,4	0,0	0,6
9	63,6	2,0	89,0	4,6	90,9	1,9	98,7	1,3	0
10	99,4	0,6	99,4	0,6	99,4	0,6	99,4	0,6	0
11	82,7	1,4	98,6	1,4	98,6	0,0	98,6	0,0	1,4
1a	63,6	0,0	94,0	4,0	93,4	2,6	94,0	2,6	3,4
2a	98,7	0,0	98,7	0,0	98,7	0,0	98,7	0,0	1,3
3a	83,7	3,3	98,7	0,0	98,7	0,0	98,7	0,0	1,3
4a	100,0	0,0	100,0	0,0	100,0	0,0	100,0	0,0	0
5a	97,6	0,0	97,6	0,0	100,0	0,0	100,0	0,0	0
6a	94,9	2,6	94,9	2,6	98,7	1,3	98,7	1,3	0
7a	80,5	5,2	93,5	5,2	96,8	0,0	97,4	0,0	2,6
8a	59,1	5,7	87,4	5,7	98,1	0,0	98,1	0,0	1,9
9a	59,5	3,2	86,1	6,3	88,6	4,4	97,5	2,5	0
10a	99,3	0,0	99,3	0,0	99,3	0,0	99,3	0,0	0,7
11a	80,4	12,3	92,0	5,8	93,5	4,3	99,3	0,0	0,7
Min	55,1	0	85	0	88,1	0	94	0	3,4
Max	100	12,3	100	11,3	100	5	100	3,5	0
μ	81,6	2,8	94,8	2,9	96,7	1,2	98,5	0,6	0,9
σ	17,0	3,5	4,7	3,2	3,7	1,8	1,6	1,1	1,0

TABLE 4.2: Results of the algorithm vs ground truth for each different stage.

with an average number of merges of 1.3 among these cases.

The stage that best reduces the number of FP is the track refinement technique, decreasing it on average from 2,9% to 1,2% and for the worst case from 11,3% to 5% while improving similarly the TP rate (94,8% to 96,7%).

The track refinement based on the 3D marker segment model has a positive impact on the quality of tracks (no decrease in TP, no increase in FP). Fig. 4.18 illustrates several frames of sequence 3 and the automatically detected and labeled markers.

4.5 3D stent Reconstructions

We have reconstructed the stents both based on the ground truth and on the automated marker detection. No visual differences could be observed between the two volumes. Therefore the results produced by our automated marker detection are fully satisfying for our application. We illustrate some of the 3D stent reconstructions hereafter. Fig.




FIGURE 4.18: Automated marker detection and labeling on sequence 3.

4.19 (left) illustrates the 3D stent reconstruction of sequence 3 reconstructed with the state of the art similarity transform.




FIGURE 4.19: Left: 3D stent model reconstructed from sequence 3. Right: 3D stent of sequence 6 (bifurcation case)

In this case, the 3D model allows to visualize the overall shape of the stent and the details of the stent struts. Axial slices allow to confirm the circular shape of the stent cross sections and its harmonious deployment. Fig. 4.19 (right) presents a 3D reconstruction where 2 stents are used simultaneously for a bifurcation case. Fig. 4.20 illustrates 3 additional cases. Note that the image at the center is the same case as Fig. 4.19(left) without balloon inflation. We observe that the quality of the visualization is lower. This observation is general (balloon inflated reconstructed stents are of better quality) and this is a consequence of the fact that the inflated balloon stabilizes the landmarks with respect to the stent. Moreover when the balloon is deflated the guide-wire often sticks to the stent (instead of being at the center of the vessel lumen) and impairs the precise stent visualization.




FIGURE 4.20: 3D stent reconstructions from left to right: sequence 2, 3a and 1a.

On Fig. 4.21 we can observe axial and longitudinal cross sections of the un-optimally deployed stent presented in Fig. 4.2. It allows to quantify the residual stenosis.




FIGURE 4.21: Left longitudinal slice / Right axial slice of the reconstructed stent (image from [Per08]).

4.5.1 Guide-wire subtraction and registration

The guide-wire subtraction is visually satisfying on the projection images as illustrated previously in Fig. 4.14. In the 3D reconstruction however a mark persists, though much weaker than the original guide-wire signal. Note on Fig. 4.22 that the stent border that was hidden by the guide-wire is restored by the guide-wire subtraction technique.




FIGURE 4.22: 3D stent slice: Left classical processing, right guide-wire subtracted between the markers. The stent border hidden by the guide-wire on the left is visible on the right, as indicated by the white arrows. The guide-wire subtraction process attenuates the guide-wire signal, but in the 3D volume a mark still persists.

We did not perform a quantitative comparison of the different registration techniques. However, we present here some preliminary qualitative results that demonstrate the interest of the tools we developed on clinical cases. We first note, that the impact of the registration technique on the reconstructed volumes is visible on the longer stents only. For the small ones, all the registrations yield visually equivalent results. This behavior

was expected, and is similar to the 2D case (see Section 2.4.3.1). We illustrate in Fig. 4.23 and 4.24 a long stent case (20mm) where the quality of the results can be ranked in the following order:

$$\text{similarity} < \text{radius preserving affine} < \text{non-linear} \quad (4.17)$$

In this case, the combination of guide-wire subtraction and the non-linear registration yields a significant improvement: the stent is sharper and its shape better defined, both in volume rendering and in the slices of the volume.




FIGURE 4.23: Left: state of the art. The transform to register the markers is a similarity and the guide-wire is not subtracted; center: radius preserving affine registration and guide-wire subtraction; right: non-linear registration and guide-wire subtraction. The three rows illustrate different viewing angles of the same stent. Observe on top right that the strut pattern appears in the upper side of the stent. This pattern is not visible on the left image. On the bottom right the general shape of the stent can be easily assessed (this it is not the case on the left).




FIGURE 4.24: Comparison of the similarity transform (right) and of the non-linear one (left), in the slices of a reconstructed stent volume (same example as Fig. 4.23). Guide-wire has been subtracted in both cases.

4.6 Discussion

4.6.1 Comparison with state of the art

Besides our team, three other ones have attempted 3D coronary stent reconstruction in the cathlab. The one from Cedars Sinai [CMW97] led the first study on this topic in '97 that demonstrated the feasibility of the approach. However they did not follow this research path any further.

The second one, from Philips Healthcare, started more recently to tackle this topic and produced numerous publications [MSG⁺06, MB06, SMG07, SFL⁺08, SF09, SFL⁺09a] including human studies. This is the *only* other team who published on automatic marker detection and tracking in spin acquisitions. The general architecture of their

tracking algorithm shares some similarities with ours. First they detect marker candidates, then form couples, and build tracks based on temporal constraints. The detected markers are reconstructed in 3D, and the 3D model is used to further enhance the 2D track. The results they report, in their latest publication [SFL⁺09a] on automated marker tracking are: TP 84%, FN 18% and FP 2% with the necessity of user interaction in 1 case out of the 10 cases. These results are in the same order of magnitude as ours, but inferior. However, the databases are different and several external factors, like the acquisition protocol, the image quality, the devices (stents and balloons) that are used, can impact these success rates. Regarding motion compensation, they register the projected 3D marker model to the actual detected markers with a linear registration. We proposed two variations over this concept: a modification to the linear registration to better preserve the stent radius, and a non-linear registration that can account for more complex motions (like vessel bending).

The third team [RLH11b] relies on manually detected markers and focuses on the 3D reconstruction technique. They first reconstruct the trajectory of the markers in 3D (yielding a 3D+t marker model) that they use for the reconstruction of the stent volume. We have not investigated this research axis on our side.

Finally, we note that there is no previous work about guide-wire subtraction that we could compare our method with.

4.6.2 On the acquisition

Two points can be optimized to increase the success rate of the automated 3D stent reconstruction. On the first hand the position of the axis of the spin with respect to the axis of the balloon. On the second hand the low pressure inflation of the balloon to minimize the relative motion between the balloon and the stent.

If the axis of the spin is chosen to be parallel to the one of the stent, the stent will appear with a rather constant length. Conversely, if the axes are orthogonal, the length of the stent will vary along the spin acquisition passing by a singular case where it is null. In the frames close to this singular point the markers appear as one single point. This situation is clearly less favorable for a tracking algorithm since it makes the balloon appearance less invariant and the labeling less straightforward. Moreover the estimation of the registration parameters becomes singular. Only the translation parameter can be estimated, while the rotation and scaling are degenerated. Error propagation shows that an uncertainty ϵ in the position of the exact marker center generates an uncertainty in the scaling factor and on the cosine of the rotation angle proportional to ϵ/l , where l is the length of the marker segment. As a result there is a higher uncertainty on the transform in the angles where the stent projection is shorter. Fig. 4.25 illustrates the inverse of the length of the marker segment and the estimated rotation angle as a

function of the frame index. We observe that the rotation angle becomes unstable when the length of the marker segment decreases. We finally want to recall that when the stent length decreases as a result of foreshortening, the stent diameter stays constant. Therefore the uncertainty on the transformation parameters impacts the diameter of the stent.

It has been observed on the reconstructed stents in the Results section 4.4 that balloon inflation presents two advantages: it enhances the quality of the reconstructed stent, and positions the guide-wire further from the stent border (see Fig 4.26). We also observed less difference between the non-linear and linear registrations on the balloon inflated cases. This is probably due to the fact that the inflated balloon prevents the artery from bending. Therefore balloon inflation reduces the need for sophisticated transforms and for guide-wire subtraction. However it is an additional constraint on the work-flow that may severely reduce its acceptability in daily practice.




FIGURE 4.25: Stability of the estimation of the parameters of the 2D registration versus the marker segment length. Left: inverse of the marker segment length (in pixels) vs the frame index. Right: estimated rotation angle in radians vs the frame index.




FIGURE 4.26: Volume rendering of a same stent, with (left) and without (right) balloon inflation. Observe that the inflation of the balloon is (i) beneficial to the quality of the enhanced stent and that (ii) prevents the guide-wire from overlapping with the stent border. This illustration case is the one where the impact of balloon inflation is the most striking in our database.

4.6.3 Future work

The overall results of the automated marker detection and labeling algorithm (TP 98%) are very close to the maximal achievable result. Since we noticed no difference between the reconstruction based on automatically detected and manually segmented markers, enhancing TP can only have marginal impact on the the 3D stent. However it can be useful for other applications, like enhanced stabilized display [SF09].

Developments that are relevant for our application are the one that can further enhance the quality of the reconstructed stents. Taking into account a 3D+t guide-wire models can be a future research axis.

Ultimately, compensating the motion of the stent versus the landmarks may enable to yield similar results with and without balloon inflation and would simplify the workflow. Finally, a clinical validation with IVUS on a large database of cases would enable to quantify the reliability of the reconstructed stents.

4.7 Conclusion

We have proposed an algorithm for the fully automated detection and labeling of the balloon marker-balls in rotational acquisition image sequences. Our bottom-up approach is building primitives of increasing level of abstraction and confidence. It exploits the 2D and 3D properties of the object to detect. It has been applied on 22 clinical sequences with a very satisfying success rate (on average 98%), the best reported so far in this context. The results of the automated detection have been used for motion compensated 3D stent reconstruction. The comparison of the reconstructed volumes with the ones reconstructed using a full manual segmentation exhibit no noticeable difference. The quality of the automated detection enables the daily use of the 3D stent reconstruction from rotational acquisitions from a workflow point of view. Moreover we showed the feasibility of guide-wire subtraction in the 3D volume by processing of the 2D projection images. We investigated the impact of several 2D registrations for motion compensation. It resulted on our dataset that there is no noticeable difference in the reconstructed between the different reconstruction. We thus recommend the similarity transform that is the simplest one from a computational standpoint. These results need now to be validated from a clinical point of view. It would be especially interesting to compare this technique to other reference ones such as IVUS.

Chapter 5

Curvilinear structure enhancement

5.1 Abstract

Curvilinear structures arise naturally in biology and medical imaging, where they are particularly abundant. Given their specific spatial representation, they require dedicated processing and have received a lot of attention from the image processing community. We first propose in this chapter¹ a review and a classification of curvilinear structure enhancement methods. Then the methodologies to compare them are reviewed. We propose a new structure to process curvilinear structures, that we call the the polygonal path image. We derive from \mathfrak{P} some linear structure enhancement and analysis algorithms. We show that \mathfrak{P} has several interesting properties: it generalizes the concepts of several methods into a unified formalism, and has very simple and intuitive parametrization, and enables the control of the smoothness and length of the structures to detect. Moreover it benefits from an efficient computational scheme. In order to evaluate the applicability of the method, we estimate quantitatively its performance within the ROC curve formalism on a practical example, the detection of guide-wires in X-ray images. The results outperform previous state of the art techniques.

5.2 Introduction

Curvilinear structures appear naturally in the human body and thus in medical images. Their segmentation, that enables various diagnostic and interventional application, is an extensively studied topic [MJS⁺04, MJSU03, HTD05, OG10, RC08, WZL⁺07, Mei10, LABFL09, FRvAtHR06, PWBB97, BVN03, FRvAtHR06, SKUF07, BVG09, BAG⁺07,

¹I want to thank Hugues Talbot for his supervision on this study, Olena Tankyevych for her collaboration on the comparison of line enhancement techniques and Nicolas Honnorat for the annotated ground truth on the clinical image database.

GBK⁺09, AGZ⁺08, HGG⁺09, WCZ⁺09, BLHS09, LBL⁺09, HVP10a, HVP10b, BWL⁺10, Lia10, BLSH10]. We propose in this chapter a new powerful technique for the enhancement of curvilinear structures in an image, based on local shortest paths. The proposed technique associates: (i) a simple formalism, (ii) the resolution of several limitations of previous state of the art techniques and (iii) efficient implementation. It only involves a very reduced number of parameters (typically 2) that have simple interpretation and, in contrast to most of the shortest paths approaches, it does not require either a starting or an ending point. One of the main advantages of our proposed formalism is that it unifies local, semi-local, and global techniques in a single framework where one can set the scale of analysis. We first review state of the art techniques for curvilinear structure enhancement and methodologies to compare them. Then we expose a theoretical description of our method and discuss its links with other known techniques (namely local paths [Vin98, TA07], geodesic voting [RC08] and Hough transform [Pau59]). We finally quantify its performance for the difficult task of guide-wire detection in X-ray fluoroscopy and compare it to other state of the art techniques.

5.3 Background

The segmentation of curvilinear structures in medical imaging is an extensively studied topic that covers a wide variety of naturally elongated biological structures and medical tools: neurites [MJS⁺04, MJSU03], microtubules [HTD05, OG10], microglia [RC08], neurones/axons [WZL⁺07, Mei10], ducts in mammograms [ZABT04], cerebral structures [MvdEV96], vessels [LABFL09], interventional tools like catheters and guide-wires [PWBB97, BVN03, FRvAtHR06, SKUF07, BVG09, BAG⁺07, GBK⁺09, AGZ⁺08, HGG⁺09, WCZ⁺09, BLHS09, LBL⁺09, HVP10a, HVP10b, BWL⁺10, Lia10, BLSH10] etc. We therefore start this section with a description of the state of the art in curvilinear structure segmentation. Then, we present a challenging application in the context of interventional radiology: the segmentation of guide-wires in X-ray images. Contrarily to the approach presented in Chapter 2, where we relied on the presence of contrasted markerballs to constrain guide-wire detection, we aim here at segmenting it without any external information. Finally we explain the methodology we used to compare curvilinear structure detectors.

5.3.1 State of the art

Authors generally describe line detection techniques in two steps: a local pre-processing followed by a global segmentation approach [BVG09, HVP10a, HVP10b]. Looking in more details at the related literature, it appears more accurate to us to describe them in three steps (the second one being optional) that correspond to three different scales of analysis:

1. Local scale: This step consists in building a map (sometimes called feature map) that represents the probability of presence of an elongated structure at each pixel and possibly its orientation [FNVV98, SNA⁺97, Lin98, PWBB97, LBL⁺09, OG10, HTD05, MJS⁺04, WZL⁺07, BLHS09, MvdEV96, KS05, BVG09, FAoTMLVG91, JU04, BAG⁺07, WCZ⁺09, HVP10a, HVP10b, Vin98, TA07]. The computation of such features generally involves considering pixels at a short distance from the pixel to process and assumes that the structure is straight at the considered scale.
2. Semi local scale: The feature map is re-enforced by a process that tends to enhance aligned responses [FRvAtHR06, PWBB97, GM91]. Tensor voting and coherence enhancing diffusion are examples of such techniques.
3. Global scale: At this stage, simple operators like thresholding are generally not sufficient to segment the curvilinear structure. A higher level process is invoked to perform segmentation of the feature map. It either involves grouping [SB93, BAG⁺07, HVP10a] or tracking [WZL⁺07, BLHS09] (usually within a shortest path formalism). Heuristic information about the nature of the structure to segment can be embedded in the algorithm at this stage [BAG⁺07]. When several images of a slowly moving structure are segmented in a row, this stage can be performed by tracking a snake spline [BVN03, HGG⁺09].

Step 1 very often relies on the computation of the Second Order Derivatives of the image (SOD), known to characterize locally elongated structures [FNVV98, Lin98] as valleys or ridges. In the simplest cases, the feature map is obtained with a Laplacian [PWBB97]. However, most of the time, the feature map is derived from the eigenvalues/vector of the Hessian [LBL⁺09, OG10, HTD05, MJS⁺04, WZL⁺07, BLHS09, MvdEV96], that characterize the presence/absence of an elongated structure, its contrast and orientation [FNVV98, SNA⁺97, Lin98]. Some authors convolve the image with a bank of filters (each filter being dedicated to a specific orientation in the image) [KS05, BVG09] and extract from the bank of responses the local direction and intensity of the potential elongated structure. This approach can be elegantly embedded in the steerable filter formalism that reduces the amount of required computation [FAoTMLVG91, JU04] in some particular cases. It can be considered as a generalization of the Hessian since the Hessian is one of them. Finally, the learning formalism has been used to build the feature map [BAG⁺07, WCZ⁺09, HVP10a, HVP10b]. The resulting feature map can combine a large number of curvilinear structure detectors (for instance more than a dozen in [HVP10b]). An exception to these approaches is the case of local path processing [Vin98, TA07] that characterizes the local structure as following a path in the image. This notion is more general and flexible than the ones expressed by the Hessian and filter bank formalisms.

Step 2 consists in enhancing the feature map responses that are close and aligned. This enhancement of the feature map can be a filtering of the feature map, for instance when it is performed by tensor voting [FRvAtHR06]. In some other cases it involves an explicit parametric representation of the contour [PWBB97, GM91]. Generally speaking, features derived from the Hessian, are considered as first order models of the curvilinear structures (they provide its contrast and orientation or tangent line). The parametric local fits extend this local knowledge beyond the first order with parabolic [PWBB97] or circular models [GM91] that are usually valid on a larger scale.

The third and last step involves grouping the significant responses in the feature map together into full curvilinear structures. It is the conversion from a pixel-wise, or short range information to a higher level primitive. The formalism of grouping [SB93] has been successfully applied to this problem in its different variations, from the simplest to the most sophisticated:

- greedy, grouping first the responses that match best.
- hierarchical [BAG⁺07], building primitives of increasing length.
- yielding the solution of a minimization problem [HVP10a].

The most common alternate approach to grouping is tracking: starting from a point on the structure (potentially given by the user in a semi-automatic framework), the extraction of the curvilinear structures propagates the information by selecting the relevant neighbors [LBL⁺09, PWBB97]. It is very often expressed in a shortest path formalism where the feature map is used as a potential [WZL⁺07, BLHS09]. In case when the feature map contains orientation *and* contrast information, both can be embedded in the potential [OG10, HTD05, MJS⁺04]. In the more restricted number of applications where a set of images are segmented successively, this last step can be performed by snake spline tracking [BVN03, SKUF07, HGG⁺09].

5.3.2 Guide-wire detection in X-ray fluoroscopy

We are interested here in one particular application of curvilinear structure segmentation: the segmentation of guide-wires in X-ray fluoroscopy during PCI. We recall the clinical challenges that motivate this application and the technical specificities encountered in this context.

5.3.2.1 Motivations

Interventional radiology/cardiology therapies imply inserting guide-wires in the vascular system of the patients under the monitoring of low dose X-ray images, called fluoroscopic images. The field of application of such procedures is constantly growing, far beyond

the scope of angioplasty dealt in the previous chapters of this thesis, extending from aneurysm treatment to tumor embolization. Over the past years, guide-wire detection in fluoroscopic images has gained interest and maturity among the image processing researcher community, addressing a trade off between fast execution and high detection performance. Guide-wire detection enables a wide range of applications such as visualization enhancement, 3D guide-wire reconstruction and respiratory motion tracking. The automatic segmentation of the guide-wires in a wide variety of clinical conditions differs from the problem that we treated in Chapter 2. In this chapter, guide-wire detection was considered in the context of DSE where we supposed that the two radio-opaque balloon markerballs were present. The detection of these markerballs considerably simplified the task of segmenting the guide-wire. Here, we do not assume that markerballs are present. In this context, guide-wire segmentation is a particularly difficult task. We detail hereafter the related technical challenges.

5.3.2.2 The technical challenge

In X-ray images, guide-wires appear as thin, dark curves. Therefore, their detection can be achieved with methods commonly used for thin objects detection. More specifically, they are hundred of pixels long structures of typical radius 0.036mm, which translates into 1 to 5 pixels on Xray images in the cathlab. They are thus extremely anisotropic structures. A small number of guide-wires simultaneously appear in images: in most of cases a single one, rarely two to three. Their appearance obeys to the physics of the formation of X-ray images: their Contrast to Noise Ratio (CNR) is decreasing with decreasing background intensity. They are generally formed of two distinct parts: a long guide-wire body of low CNR (between 0 and 3) and a short guide-wire tip, or extremity, of higher CNR. The challenges in guide-wire segmentation arise from the low CNR of their body part and of the disturbing clutter and anatomical structures that superimpose to them. It is thus a relevant technical challenge to benchmark line enhancement techniques according to their ability to handle high noise and clutter.

5.3.3 Evaluating and comparing line enhancement methods

For evaluation purposes, we consider line enhancement as a two class detection problem where the line enhancement algorithm is the classifier: it assigns a value to each pixel that increases as the pixel is likely to belong to a curvilinear structure. We want to characterize the ability of this processing to assign larger values to the pixels of the structure to detect than to the other ones. This is classically captured by a ROC analysis that we briefly recall hereafter. Then we explain precisely how this formalism can be adapted to our guide-wire detection application.

5.3.3.1 About ROC curves

ROC stands for Receiver Operating Characteristic. It is a common tool to study the performance of a classifier and to compare several of them. It was first introduced in signal detection theory [Swe96] and later conquered a wide range of fields including medicine [Swe88] and machine learning [Bra97] where it is extremely popular nowadays. We recall here briefly the necessary background knowledge to support our use of this tool in this chapter ². We explain in this section, what is the ROC curve, what are its principal properties, and how to use it to compare classifiers.

		actual value	
		p	n
classifier output	p'	true positive TP	false positive FP
	n'	false negative FN	true negative TN
total		P	N

FIGURE 5.1: Confusion matrix

Let us consider first the characterization of a two class binary classifier b over the dataset D . Each data point $i \in D$ belongs to one of the two classes $\{p, n\}$ and is assigned a class $b(i)$ by the classifier in $\{p', n'\}$. The ideal classifier shall output p' (resp. n') for every data point in p (resp. n). The outcome of a non ideal classifier falls into one of the four possible categories summarized in the confusion matrix (Fig. 5.1): a data point that is classified p' and that actually belongs to p (resp. n) is called a true (resp. false) positive. True and false negatives are defined in a similar manner. The capital initials TP, FP, TN and FN denote here the total number of data points in each case and P (resp. N) denotes the number of points that belong to p (resp. n) in the dataset D . These quantities form the basis of the most common metrics. Those necessary to define the ROC space are the True Positive Rate (TPR) and the False Positive Rate (FPR) defined in the following equations:

²A very comprehensive introduction to ROC curves can be found in [Faw06].

$$\text{TPR} = \frac{\text{TP}}{\text{P}} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (5.1)$$

$$\text{FPR} = \frac{\text{FP}}{\text{N}} = \frac{\text{FP}}{\text{FP} + \text{TN}} \quad (5.2)$$

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (5.3)$$

$$\text{accuracy} = \frac{\text{TP} + \text{TN}}{\text{P} + \text{N}} \quad (5.4)$$

The ROC analysis is the characterization of a classifier by the couple (TPR, FPR). It is usually represented in a two axis graph (see Fig. 5.2). Several points are of particular interest in this space : The point (0,0) illustrates an extremely conservative classifier that never outputs any positives and (1,1) an extremely liberal one that always outputs positives. The point (0,1) corresponds to the ideal classifier that never commits any classification error. Generally speaking a point in ROC space is better than another if it is to the northwest (TPR is higher, FPR is lower) of the first. On Fig. 5.2, point C' thus represent a better classifier than A. The points that lie on the $y = x$ diagonal line are random choice classifiers (plotted with a dashed red line on Fig. 5.2). They have no information about the class. The points below the $y = x$ line represents classifiers that perform worse than random guessing. This triangle is therefore usually empty in ROC graphs. In such a case, simply negating the classifier (inverting its classification decisions) turns it into relevant one. The ROC characterization of a classifier has the important property of being insensitive to a change in class distribution: no matter if a given class is more represented than the other in the dataset, the ROC point of a classifier stays the same³. This is important for our application where we have far more negatives than positives in an image.

We have seen that a binary classifier generates a single ROC point. However, the line enhancement techniques we deal with are not binary classifiers: they output a continuous quantity $f(i)$, that has the only property, that a higher score indicates a higher probability of belonging to the linear structure class. A set of thresholds can be applied to this output to derive a family of binary classifiers $\{f_t\}$ according to Eq. 5.5. Each of these binary classifiers generates a point in the ROC space, so that f generates a whole set of points that are aligned on a curve called the ROC curve. An illustration of a ROC curve is shown on Fig. 5.3. Conceptually the thresholds range from $-\infty$ to $+\infty$. For the thresholds higher than $\max_{i \in D} f(i)$ the ROC point is (0,0) and for those lower than $\min_{i \in D} f(i)$ the point is (1,1). The curve is monotonous since a lower threshold produces necessarily higher FPR and TPR. From a practical standpoint, the ROC curve can be efficiently computed in one pass over the sorted values of f over D . Similarly

³This is not the case of metrics like the precision that are not normalized by the number of elements in the two classes




FIGURE 5.2: ROC space (image from [Ind09]). Point B illustrate a random choice classifier. C' is a better classifier than A. C is the negation of C'.

to ROC point comparison, the classifier with the highest ROC curve is the best. The optimal classifier is represented by a step function that is 0 for $x = 0$ and 1 otherwise.

$$f_t(i) = \begin{cases} N & \text{if } i < t \\ P & \text{if } i \geq t \end{cases} \quad (5.5)$$

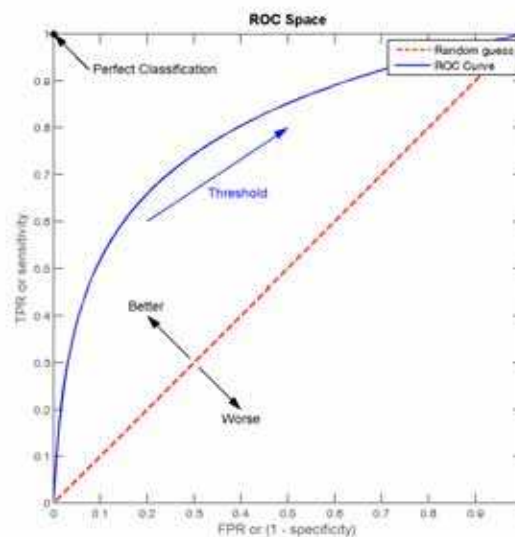


FIGURE 5.3: ROC curve illustration (image from [Ind09]).

In the general case, when a set of classifiers are compared, none of them presents a ROC curve higher than all the other ones. The curves intercept, and the highest one depends on the considered interval. However, the ROC curve can be used to derive a statistically meaningful metric to evaluate classifiers: the Area Under the Curve (AUC).

It has been shown that it is an estimate of the probability that the classifier will rank a randomly chosen positive instance higher than a randomly chosen negative instance [Faw06]. The minimal value of this metric is 0.5, for a random choice classifier, and the maximal value is 1 for a perfect classifier. The relevance of the AUC to compare classifiers versus other metrics like the accuracy has been demonstrated in the context of machine learning in [Bra97]. Efficient algorithms exist to compute the AUC from the sorted $f(i)$ data in one single pass [Faw06]. When the number of negatives is far larger than the number of positives, the left side of the ROC is the most relevant part [Faw06]. A common way to focus the ROC comparison on an interval of interest is to restrict the AUC computation to this interval. This metric is denoted the partial AUC. Finally when two classifiers are compared based on their AUC, the uncertainty on the AUC shall be taken into account.

5.3.3.2 Application to guide-wire detection

We have previously compared guide-wire segmentations in Chapter 5. We were in a context where the method to evaluate was actually producing a segmentation of the guide-wire as a smooth curve. In this case, the Hausdorff distance between the segmented guide-wire and the ground truth was a relevant tool to estimate the accuracy of the segmentation. We are more interested here in comparing classifiers that return a real value for each pixel. This value is all the greater as the probability of the pixel belonging to the guide-wire is high. In this context, ROC curves are a naturally suited tool.

Two recent papers presented comparisons of line enhancement techniques based on ROC analysis [ZABT04, AR07]. We have also used this methodology in two of our previous works [BVG09, Tan10] in the context of guide-wire detection. Another group of authors [BAG⁺07, HVP10a, HVP10b] that do not exactly rely on this methodology, computed metrics related to precision and TPR (but not to FPR). We examine here what are these different methodologies and present ours. Then we synthesize the results of the previous comparisons of line enhancement methods and draw guidelines for the design of successful techniques.

The first point to address to compare methodologies is the choice of the aspect of line enhancement techniques that one wants to evaluate. All the authors have compared detection capability [ZABT04, AR07, BAG⁺07, HVP10a, HVP10b]. Some authors additionally estimated the ability to compute local orientation [ZABT04, AR07], localization error [ZABT04], execution time [BVG09, AR07], robustness to parameter setting [AR07] and to local curvature [BVG09]. Each of these aspects is associated with a dedicated methodology. This is beyond the scope of the present study that focuses on the common

aspect of all these papers, the *detection* capability of the techniques, that is also to our eyes the most important one. A review of published comparisons of line detection techniques is summarized in Table 5.1 where we indicate the methodology, the compared techniques and the outcome. The AUC is used to compare line detection techniques in [ZABT04, AR07]. In [BVG09, Tan10] the ROC curves are computed and compared qualitatively, but the quantitative comparison is performed on the TPR for a fix FPR of 1% in [BVG09] and on accuracy in [Tan10]. In [BAG⁺07, HVP10a, HVP10b] the performance evaluation is based on metrics derived from precision and TPR. Finally, in [HVP10b] several methods are compared based on the qualitative comparison of the ROC curves. We reject the qualitative comparison of ROC curves since it is not properly defined if the curves intersect. The metrics that are not based on TPR and FPR are usually not insensitive to class skew. Therefore we do not select them. The TPR at 1% false positive is not integrating the response of the line enhancement technique over a significant interval of FPR/TPR trade-off and thus lacks generality. Therefore, we propose to evaluate line detection techniques based on the **partial AUC**. It has the following properties:

- it is based on TPR and FPR,
- measures the ability of the line enhancement technique to set larger values to the pixels belonging to the linear structures
- it is focused on an interval of FPR rate.

The choice of the interval of interest is guided by the strong class skew of our problem: negatives are typically some hundred times more numerous than positives in our image. Therefore we focus on low FPR rate in the range [0, 5%]. We will also compute additional metrics derived from TPR, **precision** and **accuracy** in order to derive results similar to previous publications (see Table 5.1).

The computation of all these quantities is straightforwardly derived from TP, FP, TN and FN as explained previously. The definition of these elementary quantities is however different from one author to another in the case of curvilinear structures. In accordance to [ZABT04, HVP10a, HVP10b, Tan10] we propose to define a ground truth \mathcal{T} on each image (annotated by a operator) that is a one pixel thin line. It is the representation of the centerline, or backbone, of the structure. We denote by \mathcal{D} the set of points detected in an image. In order to take into account the real width of the curvilinear structure, we define the true positive as the pixels of \mathcal{T} that lie closer to \mathcal{D} than a pre-defined distance d , where d is typically the radius (width) of the curvilinear structure. It can be simply computed by $|\mathcal{T} \cap \delta_d(\mathcal{D})|$ where $\delta_d(A)$ is the morphological dilation of A by a

ball structuring element of radius d . Let $|E|$ denote the cardinal of set E . TPR is then defined by:

$$\text{TP} = |\mathcal{T} \cap \delta_d(\mathcal{D})| \quad (5.6)$$

$$\text{TPR} = \frac{|\mathcal{T} \cap \delta_d(\mathcal{D})|}{|\mathcal{T}|} \quad (5.7)$$

False positives are defined as the detected pixels that lie further from \mathcal{T} than a distance d . Denoting A^c the complementary of the set A in the image, the FPR rate is computed by:

$$\text{FP} = |\delta_d(\mathcal{T})^c \cap \mathcal{D}| \quad (5.8)$$

$$\text{FPR} = \frac{|\delta_d(\mathcal{T})^c \cap \mathcal{D}|}{|\delta_d(\mathcal{T})^c|} \quad (5.9)$$

These definitions allow us to compute ROC curves and partial AUC for any line enhancement technique. Moreover they allow us to compute the Missed Detection Rate (MDR) and False Detection Rate (FDR) defined in [BAG⁺07, HVP10a, HVP10b]. For sake of completeness, we express these quantities with the previously introduced notations:

$$\text{FDR} = \frac{\text{FP}}{\text{FP} + \text{TP}} = 1 - \text{precision} \quad (5.10)$$

$$\text{MDR} = \frac{\text{FN}}{\text{P}} = 1 - \text{TPR} \quad (5.11)$$

According to these metrics, a classifier is all the better as both MDR and FDR are low. In order to illustrate the classifiers performance, we show the segmentation results for the threshold yielding the minimum of MDR+FDR. In order not to corrupt our assessment of line enhancement techniques with the presence of other interventional tools or similar curvilinear structures we also compute the metrics associated to false positives in an ROI around the marked truth (see Fig. 5.4).

We want to end this section by raising a methodological point. As pointed out in [JM02], the output of most line enhancement techniques depends on a scale parameter. When this parameter is larger, a given linear structure produces a response that has larger spatial extent. Therefore the definition of FP that we presented has a size bias. In order to avoid this undesirable behavior we skeletonize the detected pixels at each threshold used for the computation of the ROC curve. Conceptually, it is equivalent to considering that we dispose of a set a classifiers that are the combination of thresholding and skeletonization for every threshold. In this case, the computation cannot be achieved




FIGURE 5.4: Illustration of the computation of TP, FP, TN and FN. Top row: original image, ground truth and dilated ground truth. Bottom row, ROI, output of our line enhancement technique and threshold and skeletonization of this output.

with the efficient algorithms presented in [Faw06]. We have to actually define a family of thresholds and to apply them to the output of the line enhancement techniques.

5.3.3.3 Output of previous line enhancement technique comparisons

An interesting pattern emerges from the data presented in Table 5.1. The Rotated Line Templates (RLT) techniques outperform other detectors in 3 of the 5 studies [BVG09, Tan10, ZABT04]. In one of them [Tan10] the line template is expressed in the morphological framework and in the two other ones in the linear filtering one [BVG09, ZABT04]. We believe that this success can be explained by the strong anisotropy that can be modeled with such templates. Second order derivative kernels have the drawback of linking the extent of the analysis in the direction of the curvilinear structure and perpendicularly to it with a single scale parameter. On the contrary line template can be tuned in length and width independently to match the image content. This property is key and shall be embedded in a successful line enhancement technique. We note, in accordance with this observation, that in one of the case where RLT are bested by other filters [AR07], the length of the RLT is very small (5 pixels) in regard to the lines to segment (115 pixels). Since the authors do not take advantage of the tuning capability of the RLT in their study it is expected that they do not obtain relevant results.

Paper	Data	Comparison method	Compared techniques				Best technique	Misc.
			SOD	RLT	SF	Other		
Zwiggelaar <i>et. al.</i> [ZABT04]	Synthetic	AUC	[KtB96, Lin96]	Line operator	[KtB96]	Rotated bins	RLT	One of the SOD filter is steerable. Study also evaluates orientation and localization performances. The SOD filters are steerable. Study also evaluates orientation accuracy, robustness to parameter setting and execution time. Study also evaluates robustness to curvature and execution time.
Ayres <i>et. al.</i> [AR07]	Synthetic	AUC	[FAoTMLVG91, KtB96]	[ZABT04]	[FAoTMLVG91, KtB96]	Gabor filters [Chu92]	Gabor [Chu92]	
Bismuth <i>et. al.</i> [BVG09]	Synthetic	TPR @ FPR=1%	[FNVV98]	[BVG09]	[JU04]		RLT except for very curved objects where it is outperformed by SOD.	
Tankyevych [Tan10]	Clinical	Accuracy	[FNVV98, FG87]	[ST98]	[JU04]	Path opening [TA07]	RLT [ST98]	
Honorat <i>et. al.</i> [HVP10b]	Clinical	ROC curve position	[FNVV98]	[BVG09]	[JU04]	Classifier	Learning based classifier	

TABLE 5.1: Comparison of the line enhancement techniques. SOD stands for Second Order Derivative, RLT for Rotated Line Template and SF for Steerable Filter. These are generic categories of filters that include a large number of variations (see references in the table for more details).

5.3.3.4 Selected line enhancement methods for comparison

We will compare our line enhancement method presented hereafter to state of the art line enhancement techniques. We chose one representant of each main family of technique:

- Rotated Line Template [BVG09] (linear filter)
- Second order Gaussian derivative [FNVV98] (Frangi's vesselness)

5.4 The polygonal path image

5.4.1 Motivations

We have seen in the previous section, that according to the experiments we have performed on line enhancement methods for the detection of the guide-wire [BVG09, Tan10], and to published studies [ZABT04], the best local detectors are simply segment matching (RLT). We have previously explained why we believe they achieve good performance. We want here to put this in regard with the properties of the guide-wires. We believe that RLT is well suited to the two main properties of the guide-wires (see Fig. 5.5).

- Guide-wires follow smooth curves in the images. They are not tortuous and can be efficiently approximated by simple shapes over large distances.
- Guide-wires have low contrast to noise ratio. In order to have a filter that produces significantly different responses on the guide-wire and on the background, it needs to integrate information over a large number of pixels to reduce the noise. As a rule of the thumb, the CNR of the response of a filter of length l is proportional to \sqrt{l} .

As previously mentioned, sophisticated methods, like steerable filters cannot be used with extremely elongated kernels. Therefore they fail to handle the low CNRs of the guide-wires properly. Local path techniques do not generally constrain the paths to be smooth curves as one would expect for the guide-wires. RLT, on the contrary, can be used to integrate pixel values over long distances (some tens of pixels) along straight lines. Now that we have understood the basic mechanics behind the success of segment matching techniques, we want to investigate how to improve on them. The developments we propose are applicable to guide-wire segmentation and to the more general task of segmenting smooth curves of potentially low CNR in images. The main limitation to the CNR improvement that segment matching can bring on a curve is imposed by the curvature radius of the structure. Beyond a given length, the segments do not accurately match the curve anymore. An ideal filter would compute the response on the whole curve describing the structure to segment (some hundreds of pixels in the case of guide-wires), thus yielding a significant improvement in CNR. This is not computationally achievable




FIGURE 5.5: X-ray fluoroscopy image. A guide-wire is illustrated. Note that it is a long smooth curve, with low contrast to noise ratio.

in the classical filter bank formalism that would imply to build a convolution kernel for each considered path in the image.

The general idea of our method can be summarized as follows. We observe that RLT selects at each pixel the line segment that best fits locally. Similarly, we propose to get rid of the line segment constraint and **select at each pixel the smooth path of given length that best fits.**

We hope that such a technique, that can integrate pixel information on a larger scale than segment matching, can detect curves even in the parts where they are locally less visible. It would also be potentially very robust to noise thanks to the integration of a large number of pixels along the paths. In order to constrain the smoothness of the paths we force them to be regular polygonal lines. Such an approach can be implemented efficiently by combining segment matching and local path algorithms. We call this technique “the polygonal path image” and associate it to the notation \mathfrak{P} . \mathfrak{P} , that is described in the next section, is the image analysis tool from which we propose a series of curvilinear structure processing techniques. It is defined by two parameters: the total arclength of the path and the length of each line segment forming the polygonal line. The first one controls the analysis scale and the second one the path smoothness. The first parameter allows to embed the concepts used in the three typical analysis scales (local, semi-local and global) into one single formalism. The smoothness parameter is to be set according to the *a priori* smoothness of the structures to segment.

We present \mathfrak{P} in Sections 5.4.2 and 5.4.3. An efficient algorithm to compute it is described in Section 5.4.4. We define several operators on \mathfrak{P} dedicated to curvilinear structures in Sections 5.4.5, 5.4.6 and 5.4.7. Section 5.4.8 deals with the theoretical

links with existing techniques. Quantitative results based on the ROC formalism and comparison to other techniques are presented in Section 5.5.

5.4.2 The polygonal path image: \mathfrak{P}

The image analysis tool that we propose to build is a very rich structure that contains for each pixel the locally optimal smooth path and its cost. It is a generalization around the concept of locally optimal paths [Vin98] that we first recall. Then we motivate the need for a new tool, and describe our solution.

5.4.2.1 Locally optimal paths

Interested in the segmentation of low contrast curvilinear structures, Vincent [Vin98] proposed the concept of locally optimal paths and an algorithm to compute them. The optimality of the path is defined by the minimization of a cost function. Similarly to classical shortest path techniques [CK97], it is based on a potential image I that encodes the local cost information $I(p)$ at each pixel p . The cost $J(\mathcal{P})$ of a path \mathcal{P} is defined as the sum of the values of the pixels in I along the path:

$$J(\mathcal{P}) = \sum_{p \in \mathcal{P}} I(p) \cdot \mathbf{1}_{\mathcal{P}}(p) \quad (5.12)$$

Classical shortest paths approaches require a start and an end point, between which a *globally* optimal path is computed. Vincent considers a path length, denoted L to define the notion of *locality*. Let Φ_p^L be the set of the paths of length L starting at pixel p . The locally optimal path, denoted \mathcal{P}_p , is the one that yields the minimal cost in Φ_p^L :

$$\mathcal{P}_p = \operatorname{argmin}_{\phi \in \Phi_p^L} J(\phi) \quad (5.13)$$

In order to prevent pathological situations with self intersecting paths, and to put the smoothness of the paths under control, he refines the concept and considers to:

Compute for each pixel p of an image, the minimal cost $J(p)$ for all the paths \mathcal{P} of a given length L originating in p and whose possible orientation and straightness are within a given range.

In this approach a path is thus controlled by two additional terms: orientation and straightness. The orientation is a reference angle θ and the straightness is an angular range $\delta\theta$. The three constraints, orientation, straightness and length define a search cone in which Vincent imposes that the path shall stay at every pixel. It is illustrated in Fig. 5.6. The straightness and orientation constraints are recursive: at every pixel of the path, the remaining path must fulfill them (see Fig. 5.7). Thanks to these

constraints pathological situations are avoided: paths cannot self intersect or describe a circle around the center pixel. The cost of any path of length L that is inside the angular range $[\theta - \delta\theta, \theta + \delta\theta]$ is computed and the path yielding the lowest cost is selected. In order to simplify the problem only a small number of orientations $(\theta_i)_i$, generally regularly sampled between 0 and 2π , are considered (see Fig. 5.8). The parameters $(\theta_i)_i$ and $\delta\theta$ are chosen so that all the orientations are within at least one cone. In some applications [TA07], the parametrization is such that the cones overlap.

With a brute force approach, the number of paths to evaluate at each pixel is extremely high. However Vincent [Vin98] proposed an efficient iterative algorithm for this task (discussed in Section 5.4.4).




FIGURE 5.6: Illustration of the definition of the paths by Vincent. The center pixel, where the path originates, is p . l is the length of the path, θ_i its orientation, and the angular range $\delta\theta$ the straightness. The search cone in which paths are constrained to be located is shaded on the figure.




FIGURE 5.7: Illustration of the recursivity of the cone constraint. A path originating from a pixel p shall stay in the cone placed in any of its pixels p' . The path of (a) verifies this property illustrated in (b). The path of (c) does not verify this property (see (d)) and will thus not be considered by the technique.




FIGURE 5.8: Triangular search cones defined by Vincent. These four cones of straightness $\pi/2$ cover the whole angular range.

This approach has been successfully extended to define efficient algorithms that perform morphological operations on paths [TA07]. One of the limitations of the straightforward implementation of the minimal local path [Vin98] or the path opening [TA07]

in our context is that, in spite of the cone constraint, the paths are still allowed to be tortuous. It results in paths that can wander in the image, attracted by local noise fluctuations to minimize their cost. Fig. 5.9 shows the example of such a minimal path. The second limitation is that in his article Vincent considers to store the cost of each optimal path, and not the path itself. However, the cost of the path carries very little information compared to the path about the geometry of the curvilinear structures present in the image.




FIGURE 5.9: Left a curvilinear structure in a very noisy environment. Right, minimal path computed between the top and the bottom of the image (illustration from [Vin98]). It is extremely tortuous, and does not follow the smooth curves that are naturally present in medical images.

In order to remove the first limitation we propose a simple way to control the smoothness of the paths: we force them to be polygonal. The tortuosity is controlled by the length of the edges forming the polygonal path⁴. If this length is one pixel we are back to the original formalism. With increasing edge length, the number of allowed directional changes along the path is decreased. Similarly to [Vin98] and [TA07] we divide the neighborhood of each pixel into cones of orientation θ_i , and we constrain each path to be inside the search cone at each of its vertices. The cone constraint allows to further control the smoothness of the paths, to avoid pathological situations and enables the use of efficient algorithms. Finally we additionally constrain the polygon to be formed of segments of equal length. Let l be the length of each segment of the polygon, n the number of segments and $L = nl$ the total arclength of each path. We denote $\Phi_{p,l,\theta}^L$ the set of regular polygonal paths of segment l , total length L , inside the cone of orientation θ and originating at pixel p . The locally optimal path selected by the polygonal path

⁴The discrete line segments are computed with the Bresenham algorithm.

image is:

$$\mathcal{P}_p = \underset{\theta \in (\theta_i)_i}{\operatorname{argmin}} \underset{\phi \in \Phi_{p,l,\theta}^L}{\operatorname{argmin}} J(\phi) \quad (5.14)$$

We note that contrarily to most local path techniques that only compute the cost of the selected path, we compute and store the whole paths. We illustrate in the next section the local geometrical information that is encoded in the paths. We further derive a set of curvilinear structure processing techniques, that rely on the use of the whole paths on not on their cost.

Besides, the polygonal path image has the nice property of only depending on two parameters : the length of the base line segment and the number of line segments in the path. These two parameters can be easily set. The length of the base segment shall be the largest possible provided that it fits in the curve to detect and the number of line segments shall be longest possible so that the curve stays in the cone. They are directly related to the typical shapes of guide-wires on X-ray images and can be adapted to the application. For instance in cardiology where arteries are less tortuous than in neurology, they can be set to larger values. From a higher level perspective the product of these two parameters is the arclength of the paths to consider that is equivalent of the analysis scale presented in Section 5.3.1 and the length defines the smoothness of the paths to analyze. With one single segment in the polygon we compute the equivalent of what we denoted the local processing. With a few segments a semi-local processing similar to those that tend to re-inforce aligned responses of the local processing. With a large number of segments (typically ten) the procedure builds paths of length similar to the one obtained with a grouping or a tracking technique applied to a local processing. The main differences are that one path is built originating from each pixel and that the paths are formed of segments of equal length.

5.4.3 Structure of the path image

In order to illustrate the content of the path image and to understand its structure we have represented some paths originating from random locations on clinical and synthetic images on Fig. 5.12, 5.13 and 5.14.

We observe that the paths tend to converge into bundles around the main linear structures. From a shortest path perspective, the linear structures are similar to highways where the speed is high. The paths originating outside of the linear structures are immediately directed on the linear structures, taking the shortest possible way to reach the highway. Those that start directly on linear structures stay on them until the end or until the linear structure cannot be followed anymore because it escapes the cone. Therefore many paths have large parts in common. In the areas where there is no linear structure to attract the path, they wander randomly in noise. In the example we depict the path that we built are so long that starting from virtually anywhere in the image

they finally end up stuck to a linear structure. We can also observe some discontinuities in the path image. Indeed at a pixel that is on the equivalent of the medial axis of the linear structures for our problem, there are several shortest path of equal length that lead to different parts of the linear structures. Therefore some paths originating from pixels in the close vicinity of the medial axis will be attracted by one linear structure and others by another one. We illustrated this situation in fig 5.13.

Finally the set of paths that intersect at a pixel convey some information about the local geometry of the linear structures at the intersection point. Fig. 5.14 illustrates all the paths intersecting at a point of the guide-wire. We observe not only the high number of paths crossing at this point (nearly 2000 in this case) as previously explained but also that these paths are well aligned with the guide-wire in a short neighborhood around the considered pixel. They can thus be used to derive geometrical properties like local orientation or curvature for instance.

In order to write more specifically about the elements of the \mathfrak{P} , we recall the following notations: We call \mathcal{P}_q the path originating from the pixel q . The cost of the path \mathcal{P}_q is denoted $J(q)$, or $J(\mathcal{P}_q)$.

5.4.4 Efficient algorithm for the computation of \mathfrak{P}

In this section we explain how to build a structure that contains the minimal path of given arclength originating from each image pixel. It can be computed efficiently thanks to an algorithm commonly used in mathematical morphology [Vin98, TA07] provided that we impose a few restrictions on the paths (the cone constraint presented in Section 5.4.2). We first recall these techniques, then explain how we adapted them to force polygonal paths.

Let us describe the algorithm of Vincent [Vin98] in the case of 4 cones where $\delta\theta_i = \pi/4$, of directions $\theta_1 = 0, \theta_2 = \pi/2, \theta_3 = \pi, \theta_4 = 3\pi/2$ (Fig. 5.8). His algorithm computes the minimal cost for paths of length l with only $12l$ image shift, $8l$ point-wise minima and $4l$ additions. It is remarkable that this computation can be performed in a linear time whereas the number of paths increases exponentially with l (as $O(3^l)$). The idea of the algorithm is that the computations performed for one pixel can be partially reused for the neighboring ones. The algorithm that processes an image of potential I and returns the cost of the minimal path originating from each pixel in an image J is presented in Alg. 1.

In the case of \mathfrak{P} , one can observe that the cost of a path is the sum of the costs of its segments. A given segment being used in many potential paths, we propose to pre-compute the cost of each segment⁵ and to iteratively build polygonal paths of increasing number of segments. In order to compute the cost of every possible segment we compute

⁵A potential variation at this stage of the definition of the polygonal path image is to assign to each segment of the polygon the cost the minimal path joining its two extremities.




FIGURE 5.10: 500 paths originating at random locations are illustrated. We can observe how they all tend to converge to the linear structures of the image and especially to the guide-wire.




FIGURE 5.11: Extract of the top image. Observe the density of paths on the guide-wire.




FIGURE 5.12: Here we built a test image with a perfect circle. We illustrated again the content of the polygonal path image at 500 random locations. The behavior is similar to the one observed on the guide-wire image.

a set of cost images for each of them: $(I_{\theta_i})_{1 \leq i \leq N}$. I_{θ_i} is an image that contains the sum of the pixels along a discrete segment of length l and orientation θ_i starting at each pixel. There are typically $N = 2\pi l$ such images. For each of the θ_i we define t_{x_i} and t_{y_i} , the translation parameters from the center pixel to the extremity of the segment. We




FIGURE 5.13: We illustrated 5000 paths originating at random locations. We can observe on the line that is similar to the medial axis of the curvilinear structure for our problem, that the paths are attracted by different equidistant parts of the guide-wire. We approximately represented this line in black on the figure.




FIGURE 5.14: All the paths intersecting at one given point. This particular point belongs to the guide-wire.

Algorithm 1 Computation of the local minimal paths according to [Vin98]

Require: I

```

1: for  $\theta \in \{0, \pi/2, \pi, 3\pi/2\}$  do
2:    $I_\theta \leftarrow I$ ;
3:   for  $i = 1$  to  $l$  do
4:     for  $\alpha \in \{\theta - \pi/4, \theta, \theta + \pi/4\}$  do
5:        $J_\alpha \leftarrow I_\theta$  shifted by 1 pixel in direction  $-\alpha$ ;
6:     end for
7:      $I_\theta \leftarrow I_\theta + \min_\alpha J_\alpha$ ;
8:   end for
9:    $J \leftarrow \min_\theta I_\theta$ ;
10: end for
11: return  $J$ ;

```

illustrated these notations in Fig. 5.15. A given cone that is naturally defined by an interval of angles, can thus be defined by two integers, *cone.start* and *cone.end* between

1 and N . We define a path by a set of vertices $V = \{v_k, 1 \leq k \leq n\}$ where $v_k = (x_{v_k}, y_{v_k})$ and a cost c . Let us call $\mathfrak{P}_{cone}^{iter}$ the image that contains at each pixel the best polygonal path originating from this pixel of length $iter.l$ inside a given cone. In order to compute \mathfrak{P} , *i.e.* the structure that contains at each pixel the minimal polygonal path originating from this pixel, we perform the procedure described in algorithm 2. We illustrated some steps in Fig. 5.15, 5.16, 5.17. For our implementation we considered 8 overlapping cones of angular aperture $\frac{\pi}{2}$ at orientations $0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}, \pi, -\frac{\pi}{4}, -\frac{\pi}{2}, -\frac{3\pi}{4}$.

Algorithm 2 Computation of \mathfrak{P}

Require: (θ_i, tx_i, ty_i) with $1 \leq i \leq N$

```

1: for  $\theta \in (\theta_i)_i$  do
2:    $\mathfrak{P}_\theta^0(x, y).V \leftarrow \emptyset; 1 \leq x, y \leq ImageSize$ 
3:    $\mathfrak{P}_\theta^0(x, y).c \leftarrow 0; 1 \leq x, y \leq ImageSize$ 
4:   for  $j = 1$  to  $n$  do
5:     for  $x = 1$  to  $ImageSize$  do
6:       for  $y = 1$  to  $ImageSize$  do
7:         for  $k = cone.start$  to  $cone.end$  do
8:            $J_k \leftarrow \mathfrak{P}_\theta^{j-1}(x + tx_k, y + ty_k) + I_{\theta_k}(x, y)$ 
9:         end for
10:         $k_{opt} \leftarrow argmin_k \{J_k\}$ 
11:         $\mathfrak{P}_\theta^j(x, y).c \leftarrow J_{k_{opt}}$ 
12:         $\mathfrak{P}_\theta^j(x, y).V \leftarrow \mathfrak{P}_\theta^{j-1}(x, y).V \cup (x + tx_{k_{opt}}, y + ty_{k_{opt}})$ 
13:       end for
14:     end for
15:   end for
16: end for
17: for  $x = 1$  to  $ImageSize$  do
18:   for  $y = 1$  to  $ImageSize$  do
19:      $\theta_{opt} \leftarrow argmin_\theta \{\mathfrak{P}_\theta^N(x, y).c\}$ 
20:      $\mathfrak{P}(x, y) \leftarrow \mathfrak{P}_{\theta_{opt}}^N(x, y)$ 
21:   end for
22: end for
23: return  $\mathfrak{P}$ 

```

At this stage we have built the polygonal path image. No detection has been explicitly made so far, we just built a rich structure that contains locally minimal paths and their costs for each pixel. Its construction can be efficiently implemented thanks to the presented algorithm. The only known alternatives for the efficient computation of polygonal paths in images, are approximate. The beamlets [DHJ⁺01] build a multi-resolution pyramid a line segment and Brandt [BD99] computes the minimal necessary set of orientations, spatial positions and segment lengths to approximate line segments within a given error bound.




FIGURE 5.15: Illustration of the notations for the computation of \mathfrak{P} . Let the center pixel be of coordinates (i, j) . In the image I_{θ_k} we store the sum of the pixels along line segments of orientation θ_k . Such a line segment is illustrated on the figure with plain lines. Translation parameters (t_{x_k}, t_{y_k}) are also stored. A search cone is defined by a start angle *cone.start* and an end angle *cone.end*. All the possible locations of the end of the segments in a given search cone are illustrated by the arc of circle in dashed line. These are the locations of the pixels that will be evaluated to extend the path in the successive iterations.




FIGURE 5.16: At the second iteration of the algorithm, each pixel p_i contains the minimal path originating from it of length l . In order to compute the minimal path of 2 segments (length $2l$) originating from pixel p we consider each pixel p_i and note that the minimal path originating from p and passing through p_i is simply the already computed path originating from p_i to which we add the line segment $p_i p$.

5.4.5 Curvilinear structure enhancement techniques derived from the polygonal path image

5.4.5.1 The cost image

An existing approach to building a linear feature detector out the path image consists in associating to each pixel the cost of the path originating from it (J with our notations) [Vin98]. We have illustrated the results of this strategy on 5.19. We can observe on the perfect circle image that the cost image is quite complex. Segmenting the circle out of it is not straightforward. In the case of a clinical image 5.19 the task is even more complex. The reason for this is that the paths originating from any pixel close enough to




FIGURE 5.17: We illustrated 3 typical paths of 2 segments that have been considered by the algorithm.

the linear structure are attracted by it and thus are associated to a cost that is low, and rather close to the cost of a path starting right on the linear structure. As we build very long paths, even far away from linear structures this behavior is observed. A solution to this would be to narrow the angular aperture of the cones. Most of the path far from the linear structures could not converge following the shortest geometrical path since it would not be aligned with the linear structure inside the small aperture cones. This strategy is used in [Vin98]. However, we do not want to impose that the structure to segment stays in narrow cones on long distances⁶. The solutions that we propose rely on the fact that the polygonal path image is much richer than the simple cost image. We expose hereafter two solutions that process the polygonal path image in two opposite directions: the first one takes advantage of the redundancy of the paths to re-enforce the linear structures and the second one eliminates all the redundant paths conserving a very small set of them, most representative of the actual linear structures.

5.4.5.2 Path voting

Instead of having each pixel basically contributing to all the paths passing through it, that we typically summarize in the cost image, we propose to have each minimal path vote for all the pixels it passes through. This way, all the paths that form a bundle around the linear structures will re-enforce this structure. In the traditional approach a very contrasty pixel will spread its contrast on all the paths going through it creating a star effect. Conversely, in the proposed approach, all these paths vote for the pixels on which they intersect, making them more visible. This approach shares some similarities with the one of [Car07] who first had also observed how minimal paths are attracted by linear structures. This methodology has recently been adopted in geodesic voting

⁶This is for instance false for guide-wires.




FIGURE 5.18: On a synthetic image representing a circle we have computed \mathfrak{P} and associated to each pixel the cost of the path originating from it. We can observe that the image is complex and that linear structures tend to extend to neighbor pixels.




FIGURE 5.19: We have computed \mathfrak{P} on the clinical from Fig. 5.5 and associated to each pixel the cost of the path originating from it. We can observe that the image is complex and that linear structures tend to extend to neighbor pixels.

[RC08]. We have several options to aggregate the votes of the paths. Either each path can vote for one, which is equivalent to counting the number of paths intersecting at one pixel (we call the resulting image the "path voting image" and write it $\vartheta(\mathfrak{P})$), or each path can vote according to its cost (the resulting image is denoted $\vartheta_J(\mathfrak{P})$).

$$\vartheta(\mathfrak{P}) = \sum_{p \in I} \mathbf{1}_{\mathcal{P}_p} \quad (5.15)$$

$$\vartheta_J(\mathfrak{P}) = \sum_{p \in I} f(J(p)) \mathbf{1}_{\mathcal{P}_p} \quad (5.16)$$

Where f is a decreasing function in order to assign larger weights to the low cost paths. In practice, both technique gave quite similar results. From a qualitative standpoint $\vartheta(\mathfrak{P})$ was slightly better.

In practice path voting is a powerful tool that yields very interesting results. We use it here to illustrate the relevance of the polygonal structure of the paths (see Fig.



FIGURE 5.20: Path voting on the image illustrated on Fig. 5.5. See section 5.5.3 for details on the pre-processing.



FIGURE 5.21: A more aggressive contrast setting (saturation) of fig 5.20 shows lower values of the image. We see some path that are attracted by the guide-wire. Especially close to the tip.



FIGURE 5.22: Even more aggressive contrast setting of image fig 5.20 enables to see its complex structure.

5.26). Without the polygonal structures paths are very tortuous and in low signal level areas wander randomly through noise to match precisely the fluctuations of the image content. Polygonal regularization allows to more accurately fill the low signal level gaps.



FIGURE 5.23: Similar image as fig 5.20 except that constrained path smoothness, rejecting any path of tortuosity lower than 0.6



FIGURE 5.24: A more aggressive contrast setting (saturation) of fig 5.23 shows lower values of the image. We see some path that are attracted by the guide-wire. Especially close to the tip, however there are less such path than in 5.21 thanks to the tortuosity constraint.




FIGURE 5.25: Even more aggressive contrast setting of image fig 5.20 enables to see its complex structure. Compared to fig 5.22 we see that the most tortuous path have been removed.




FIGURE 5.26: $\vartheta(\mathfrak{P})$ with 3 pixel long polygonal segments (top) and 11 pixel long segments (bottom) for a similar total path length. In the left column (nominal windowing) we observe that the curvilinear structure appear tortuous with short path segments and smooth with longer segments (as in the original image).

5.4.5.3 Pruning

A completely opposite way to exploit the structure of the path image is to prune paths to get rid of the redundancy, retaining only a small set of relevant, non redundant minimal paths. In order to do so we define a relation that allows to determine the paths that are neighbor to a given path. We retain the path of minimal cost over the whole image and prune all the paths that are neighbors to it. They correspond to minor variations around the minimal one. The one with the best cost is likely to be the one most aligned with the actual linear structures of the image. We iterate this process in a greedy fashion selecting the path of minimal cost again and pruning its neighbors, until there is no more paths in the path image or until some stopping criterion is reached. The small set of paths finally retained (typically less than 100) describe the linear features in the image. We call them "non redundant paths". The relation defining neighbor paths can be defined in several ways. We propose to consider the following one : a path \mathcal{P}_i is neighbor to a path \mathcal{P}_{min} if their partial Hausdorff distance [HKR93] $h_K(\mathcal{P}_i, \mathcal{P}_{min})$ is below a given distance d (see Eq. 5.17).

$$h_K(\mathcal{P}_i, \mathcal{P}_{min}) = K^{th}_{p \in \mathcal{P}_i} \min_{q \in \mathcal{P}_{min}} \text{dist}(p, q) < d \quad (5.17)$$

In practice, for the sake of computational speed, we approximated this relation by counting the percentage of vertex of \mathcal{P}_i closer to \mathcal{P}_{min} than d . We used it to generate the results we present.




FIGURE 5.27: We illustrate here the result of path pruning with $x = 50\%$. We can observe that among the small number of paths selected, the ones illustrating the guide-wire are present. Further path selection, based on their cost for instance can be performed to better isolate the guide-wire.

5.4.6 Selecting path according to smoothness

It can be noted in the results produced with the proposed technique that the paths that are extracted in the parts of the image where there is no relevant linear structure are especially tortuous. Similarly some paths with very abrupt directional change occur close to high contrast objects since they are attracted by them. These paths do not meet the expected requirements of a smooth curvilinear structure like a guide-wire. Therefore we propose to remove them from the path image prior to building the detection image by pruning or voting. We define to this end a tortuosity metric τ to penalize the abrupt changes in direction. In the cones we defined the paths cannot undergo a directional change of more than 90° . Therefore we design a metric that will be 0 for a path change of 90° and 1 for a perfectly straight path. In order to do so we simply compute the product of the normalized scalar product between each two consecutive segments of the path.

Let $\mathbf{V}(k)$ be the vector formed by two consecutive vertexes of the path \mathcal{P} , $\tau(\mathcal{P})$ is given by⁷:

⁷Several other tortuosity metrics can be considered like $\min_{1 \leq k \leq n-2} \frac{\mathbf{V}(k) \cdot \mathbf{V}(k+1)}{\|\mathbf{V}(k)\| \cdot \|\mathbf{V}(k+1)\|}$, for instance that decreases less fast towards 0 when the length of the paths increases.

$$\tau(\mathcal{P}) = \prod_{k=1}^{n-2} \frac{\mathbf{V}(k) \cdot \mathbf{V}(k+1)}{\|\mathbf{V}(k)\| \cdot \|\mathbf{V}(k+1)\|} \quad (5.18)$$

In the case where all vectors have the same norm (the parameter l of \mathfrak{P}), this simplifies to :

$$\tau(\mathcal{P}) = \frac{1}{l^{2(n-2)}} \prod_{k=1}^{n-2} \mathbf{V}(k) \cdot \mathbf{V}(k+1) \quad (5.19)$$

The operation of voting with a path smoothness constraint is denoted $\vartheta_{\tau_{\min}}(\mathfrak{P})$:

$$\vartheta_{\tau_{\min}}(\mathfrak{P}) = \sum_{p \in I \text{ and } \tau(\mathcal{P}_p) > \tau_{\min}} \mathbf{1}_{\mathcal{P}_p} \quad (5.20)$$




FIGURE 5.28: We illustrate here the result of path pruning with $x = 50\%$ and tortuosity constraints (top 0.6, bottom 0.75). We can observe how the conjunction of pruning and constrained path smoothness help segment the guide-wire.

5.4.7 Estimating local orientation

We demonstrate here how \mathfrak{P} can be used to estimate the local geometrical properties of the curvilinear structure. We have already illustrated in Fig. 5.14 that when they reach a linear structure the paths tend to follow it and align with it. Therefore the set of paths intersecting on a pixel of a curvilinear structure can be used to derive its orientation (an angle in $[0, \pi[$). Outside of the curvilinear structures the paths follow the shortest path that leads to a linear structure. Therefore they convey the piece of information indicating where to find this closest linear structure (a vector in the plane). Several methods can be used to retrieve this information, the most important pitfall being the extraction of a typical orientation from a set of angles (since angle averaging is not mathematically correct). We propose a two step approach: first we extract an angular information (in $[0, \pi[$) at each pixel representing an average orientation of the paths. Then we find the direction that best aligns with paths. Let us consider a given

pixel p . We build the set of the paths $\{\mathcal{P}_i\}$ intersecting at p (*i.e.* $p \in \mathcal{P}_i \forall i$). For each of these path we compute the tangent unit vector at point p : $\{\mathbf{t}_i\}$. Then we find the angle $\alpha \in [0, \pi[$ such that the norm of the projection of the $\{\mathbf{t}_i\}$ over the line of orientation α is maximal. Let α_i be the angle between \mathbf{t}_i and the horizontal, α is solution of (see [RS89] for details):

$$\tan(2\alpha) = \frac{\sum_i \mathbf{t}_i^2 \sin(2\alpha_i)}{\sum_i \mathbf{t}_i^2 \cos(2\alpha_i)} \quad (5.21)$$

Among the two possible solutions we keep the one that yields the highest sum of square projections. This α is the orientation at pixel p . The vector indicating the direction to follow to get to the closest linear structure has orientation α and two possible directions. Since the paths are oriented (from the start point to the end point) so are the tangent vectors $\{\mathbf{t}_i\}$. We select between the two possible directions the one that yields the maximal sum of scalar products with the $\{\mathbf{t}_i\}$. On the curvilinear structure itself, this sums tends to 0 and the direction loses its meaning. The orientation, however, is still relevant.

We illustrate the result of this method in Fig. 5.29. We note that contrary to the orientation fields estimated with state of the art techniques that are non relevant outside of the curvilinear structures, ours has a precise, potentially useful, meaning. Finally, a similar methodology could be used to derive other local geometrical properties, such as curvature.




FIGURE 5.29: Extract of $\vartheta(\mathfrak{B})$ obtained from the X-ray image of a guide-wire. The arrow resulting from direction estimation give the direction of the closest curvilinear structure. Note that on the structure itself the direction of the arrow is not relevant (the linear structure can be found in both opposite directions), but the orientation is accurate.

5.4.8 Links with other existing techniques

5.4.8.1 Geodesic shortest path

We have seen in our state of the art review, that the shortest path approach is among the most popular ones for the segmentation of curvilinear structures. Our approach belongs to this wide family, but differs from traditional techniques on several aspects. The one we want to discuss here, is the link with the notion of geodesic distance. To the best of our knowledge, the link between the search of shortest paths inside cones and the minimization of a geodesic distance has not been previously investigated. We present here a first reflexion on this link, that we hope to pursue in our future work. The other differences with classical shortest path approaches will be discussed in the next section.

The classical approach of Cohen [CK97], defines the shortest path between two points A and B, by finding a geodesic for the following weighted distance:

$$d(A, B) = \min_y \int_0^L (w + I(y(s))) ds \quad (5.22)$$

where s is the arclength and L is the length of the curve. The minimum is considered over all curves $y(s)$ traced on the image domain that link the two end points, that is, $y(0) = A$ and $y(L) = B$. The constant w imposes regularity on the curve. $I > 0$ is a potential cost function computed from the image, that takes lower values near linear features. The obtained path is geodesic meaning that d is actually a distance in the mathematical sense of the term. It verifies the following set of properties:

$$\forall(x, y) \quad d(x, y) \geq 0 \quad (5.23)$$

$$\forall(x, y) \quad d(x, y) = 0 \Leftrightarrow x = y \quad (5.24)$$

$$\forall(x, y) \quad d(x, y) = d(y, x) \quad (5.25)$$

$$\forall(x, y, z) \quad d(x, y) \leq d(x, z) + d(z, y) \quad (5.26)$$

Let us examine how this translates to our local path approach.

Our technique also consists in computing a path that is minimal among a set of paths. However, not every possible path between the two extremities is considered. The cone constraint rejects tortuous ones. It can be interpreted as a bound that is set on the curvature of the paths. This property is highly desirable in some applications. In medical imaging for instance, where the paths correspond to biological structure, we often expect them to be smooth. We first observe that bounding the curvature of paths invalidates the triangular inequality necessary to the concept of distance (see Fig. 5.30). We show in the following developments that the path that are minimal within the cone constraint

are equivalent to minimizing a cost that is very similar to a distance: it verifies all the necessary mathematical properties except the triangular inequality that is only valid in a weakened, local version.




FIGURE 5.30: Counter example to the triangular inequality when path curvature is bounded. The “distance” between two points is defined as the minimal cost of all the paths joining two points, under the constraint that the maximum curvature of the paths shall not exceed a given bound. We build a particular potential cost image for this counter example, where: all the image has cost $+\infty$ except the three paths that are drawn. Let L_{AB} be the length of the path from A to B. We assign to each pixel on this path the cost $3/L_{AB}$. Consequently the cost of the whole path is 3. Similarly we assign to the pixels on the path from A to C the cost $1/L_{AC}$ and to those on the path from C to B the cost $1/L_{CB}$. The cost of each of these two paths is thus 1. Then, we note that the only finite cost path from A to B, is the one that is drawn (see Fig. (c)). Indeed, the only other finite cost path is the one going through C, but it does not verify the curvature bound constraint. Therefore $d(A, B) = 3$. Similarly $d(A, C) = d(C, B) = 1$. Finally: $d(A, B) > d(A, C) + d(C, B)$.

We first expose the properties of paths within a single cone. In a second step, we consider the combination of several cones together. For the sake of the clarity of the explanations, we consider here the paths as defined by Vincent in [Vin98].

Let us consider the simple case of shortest paths restricted to two opposite cones, for a given orientation θ (see Fig. 5.31). For any point A in the plane, we denote these two cones C_A^θ and $C_A^{-\theta}$. We call *conic-path*, a path that fulfills the cone constraint: the path shall never escape the cone starting at any of its pixels. We define the distance d_θ (called cone-distance) between to points A and B in the image, by the cost of the minimal conic-path joining A to B. If no conic-path exists, we set $d_\theta(A, B)$ to infinity.




FIGURE 5.31: The cone distance.

One can easily prove that d_θ is a distance in the mathematical sense of the term. The first two properties result from the strict positivity of the potential image. The third one and fourth ones (symmetry and triangular inequality), result from the symmetry of the cones.

In order to demonstrate symmetry, let us consider two points A and B in the image (see Fig. 5.31). If B is neither inside C_A^θ or $C_A^{-\theta}$, there is no conic-path from A to B. Therefore $d_\theta(A, B) = +\infty$. Given the symmetry of the cone definition, A is neither in C_B^θ or $C_B^{-\theta}$, and $d_\theta(B, A) = +\infty$. In this trivial case, symmetry is thus satisfied. Let us now consider the case where B is inside C_A^θ or $C_A^{-\theta}$. Without loss of generality let us consider that it is inside C_A^θ . Given the symmetry of the cones of orientation θ and $-\theta$, one observes that a conic-path from A to B in the cone C_A^θ is also a path from B to A in the cone $C_B^{-\theta}$. Therefore the distance is symmetric.

Regarding the triangular inequality, let us consider a third point C. If there is either no conic path from A to C or from C to B, the inequality is trivial. If there exists a conic-path from A to C and a conic-path from C to B, then, the demonstration (not detailed here) consists in observing that the union of these paths is a conic-path from A to B. Therefore triangular inequality holds.

The computation of the local shortest paths in our presentation is using a set of cones. We denote the set of orientations of these cones Θ . The cost function associated with the local path computation is λ defined by:

$$\lambda(A, B) = \min_{\theta \in \Theta} d_\theta(A, B) \quad (5.27)$$

Lambda is a semi-metric but not a distance: it verifies the 3 first conditions of a distance but not the triangular inequality (see Fig. 5.33 for a counter example that is similar to the one of bounded curvature paths). However a weaker local triangular inequality holds: given two points A and B, there is a non-void region $\mathcal{R}_{A,B}$ such that:

$$\forall C \in \mathcal{R}_{A,B} \lambda(A, B) \leq \lambda(A, C) + \lambda(B, C) \quad (5.28)$$

In order to define more precisely $\mathcal{R}_{A,B}$, we want to introduce at this stage the concept of inter-visibility parallelogram. To this end, we consider again the case where B is in C_A^θ . We additionally, consider a point C that is on a conic-path from A to B. C must fulfill the two following conditions:

$$C \in C_A^\theta \quad (5.29)$$

$$B \in C_C^\theta \quad (5.30)$$

Moreover, if a point C verifies these two constraints, there exists a conic path from A

to B passing through C: the union of the two line segments [A, C] and [C, B]. Therefore, the pixels that belong to a conic-path from A to B are exactly the ones that satisfy the two listed conditions. The set of points satisfying the first one is simply C_A^θ . Those satisfying the second one are the points C such that: $B \cap C_C^\theta \neq \emptyset$. This set is the dilation of B by the structuring element C^θ . It is thus equal to $C_B^{-\theta}$. We reach the following result: All the points that belong to a path from A to B belong to $C_B^{-\theta} \cap C_A^\theta$. We represented this area with a checkerboard pattern in Fig. 5.31. We denote it the inter-visibility parallelogram and refer to it with the notation $\mathcal{I}_{A,B}^\theta$. It can be defined, regardless of B being in C_A^θ or $C_A^{-\theta}$ by:

$$\mathcal{I}_{A,B}^\theta = (C_A^{-\theta} \cap C_B^\theta) \cup (C_A^\theta \cap C_B^{-\theta}) \quad (5.31)$$

The construction of $\mathcal{R}_{A,B}$ is a direct consequence of the concept of inter-visibility polygons. Let us consider the set of orientation $\Theta_{A,B} \subset \Theta$ so that $\forall \theta \in \Theta_{A,B} d_\theta(A, B) < \infty$ (i.e the set of cones for which there is at least one path from A to B inside the cone of orientation θ). One can prove that if the point C is chosen such that $\Theta_{A,C} = \Theta_{B,C} = \Theta_{A,B}$ the triangular inequality holds.

$$\mathcal{R}_{A,B} = \{C \mid \Theta_{A,C} = \Theta_{B,C} = \Theta_{A,B}\} \quad (5.32)$$

It can be determined in a more explicit way using the inter-visibility polygons:

$$\mathcal{R}_{A,B} = \bigcap_{\theta \in \Theta_{A,B}} \mathcal{I}_{A,B}^\theta \quad (5.33)$$

λ is very closely related to our approach: Let $\mathcal{C}_{A,l}$ be the square of edge l centered on A. The cost $\lambda_l(A)$ of the path of length l originating from a point A defined by Vincent in [Vin98] is simply:

$$\lambda_l(A) = \min_{B \in \mathcal{C}_{A,l}} \lambda(A, B) \quad (5.34)$$

If λ were a distance, λ_l would be a geodesic distance from A to the square surrounding it. This is not the case here. The triangular inequality has been weakened to a local version to restrict the set of paths in a way that enables to control their curvature.

5.4.8.2 Geodesic voting

Geodesic voting has recently been introduced by Rouchdy and Cohen in [RC08, RC09, RC11b, RC11a] to segment tree structures in medical images. While minimal paths are popular techniques to extract a structure between two end points, they are not easily applied with a high level of automation. The geodesic voting has been proposed as a technique to improve the automation by only requiring a starting point but no end




FIGURE 5.32: Counter example for the triangular inequality on λ . We consider the 4 cones used by Vincent [Vin98] that we group in two pairs of opposite cones, illustrated on (a) and (b). We build a particular potential cost image (b) where: all the image has cost ∞ except the three paths that are drawn. Let L_{AB} be the length of the path from A to B, we assign to each pixel on this path the cost $3/L_{AB}$. Consequently the whole path cost is 3. Similarly we assign to the pixels on the path from A to C the cost $1/L_{AC}$ and to those on the path from C to B the cost $1/L_{CB}$. The cost of each of these two paths is thus 1. Then we note that the only finite cost path from A to B is the one that is drawn that is included in cone (a) (see Fig. (c)). Therefore $\lambda(A, B) = 3$. Similarly $\lambda(A, C) = 1$. The only finite cost path from C to B is the one that is drawn, and it is in the cone (b) (see Fig. (d)). Consequently $\lambda(C, B) = 1$. Finally: $\lambda(A, B) > \lambda(A, C) + \lambda(C, B)$.




FIGURE 5.33: Illustrations of $\mathcal{R}_{A,B}$, for several configuration of A and B with the set of four non-overlapping cones presented in Fig 5.8. In this case $\mathcal{R}_{A,B}$ is the inter-visibility parallelogram.

point. In its original formulation a geodesic weighted distance U is computed as in the classical shortest path by fast marching from the starting point. Then a large set of "end points" is computed (all the points in the image, or random points etc...) from which shortest path are computed back to the starting point. Each path votes for all the pixels it passes through. The number of paths intersecting at a pixel is denoted by the authors the *geodesic density* and the process of computing it *geodesic voting*. This approach has been later extended to embed vessel radius information or to serve as shape prior to drive a snake levelset.

Remark: alternate approaches have been proposed in the past to get rid of the

end point but rely on extra parameters and/or a priori knowledge on the structure to segment, and are thus less general.

This approach has obvious similitudes with ours but also striking differences. Geodesic voting do not constrain the paths to be included in cones (this is not necessary in this numerical scheme). Therefore the paths can follow a larger variety of curves to match the image content. However we note that the paths in geodesic voting are global while ours are local. It is acceptable from an intellectual standpoint to constrain paths locally whereas globally this constraints would seem artificial and limiting. To better understand the differences between the approaches it is key to note that geodesic voting relies on a start point from which the action map is computed. This point plays a critical role and the minimal paths that vote converge to it. On the contrary our approach does not require either a starting or an ending point. The path do not all converge to a unique pixel but start from each pixel of the image. With the regular fast marching formalism this would require to repeat all the computations for each pixel. In our approach thanks to the constraint that the paths are included in the cones the computation is performed for all the pixels of the image simultaneously. Moreover, it has been shown in the geodesic voting approach that if the set of end points is not properly scattered all over the image some structures have no paths passing through them (this is referred as the shadow problem). In our approach all the pixels of the images are taken into account avoiding this pitfall. Finally if the SNR of the structure to segment is low in the potential we suspect that the geodesic path go through shortcuts to reach the starting point in the geodesic voting approach when the end points are too far from it [BC09]. In classical minimal path approaches this problem is addressed by computing smaller paths between key points [BC09] (see Fig. 5.34). This is an artificial way of making the notion of shortest path local instead of global. This problem either does not appear with our approach that builds path of fix length in a local fashion (the length of the path is *not* dependent on the position of the point in the image).




FIGURE 5.34: Images extracted from [BC09]

5.4.8.3 Hough transform

The Hough transform [Pau59, DH72] is a widely used technique for pattern recognition in images. In its most famous form it relies on a parametric representation of the shape to recognize. After discretization (sampling) of the parameter space, two approaches can be used:

- for each possible value of the parameters all the pixels of the image intersected by the pattern participate in the vote.
- the image is spanned and each pixel votes for all the parameters sets describing a pattern intersecting this pixel.

The process we use to select local optimal paths is very similar to a hough transform. For a given starting point, each pixel in the cone votes with its value in the feature map for all the curves passing through it. The curve yielding the extremal vote is selected as the local optimal path for this starting point. The main difference with Hough is that voting does not occur explicitly in the parameter space. Regarding our explicit voting scheme, it is somewhat the dual approach of the Hough transform: instead of pixels voting for curves, curves vote for pixels. In our case a set of curves (the optimal paths) are selected in the image and vote for all the pixels they pass through. Deeper links may exist between the techniques that shall be explored into future work.

5.5 Results

We present here the application of the polygonal path image to the detection of guide-wires in interventional cardiology.

5.5.1 Database

Our clinical case database is constituted of 12 clinical sequences illustrating PCI angioplasty exams from which we isolated 9 successive frames. It represents a total of 108 frames where we manually annotated the position of the guide-wire. In these images one can observe the injection catheter, the guide-wire, the angioplasty and the anatomical background. Some of them additionally display stents and sternal wires.

5.5.2 Setting of the techniques

Each technique requires some specific parameters to be set. We chose for each one the setting that appeared the most optimal to us. Regarding vesselness, we set α and β to the values proposed in Frangi's article [FNVV98]. The scale factor σ , was set to the approximate guide-wire radius: 2 pixels. For the RLT we relied on a previous study with a similar technique [Tan10] that concluded that a length of 61 pixels and a width of 3

pixels was optimal for guide-wire detection. We note that these parameters generate very anisotropic filters compared to the Gaussian derivatives used in the vesselness. For the two new techniques that we presented in this chapter $\vartheta(\mathfrak{P})$ and $\vartheta_{\tau_{\min}}(\mathfrak{P})$, we set : $l = 11$, $niter = 10$, $tortuosity = 0.6$ and applied them on half resolution images for sake of computational speed. In terms of scale, we note that with these settings, the vesselness integrates pixels along the curvilinear structure over a range of ≈ 8 pixels ($= \pm 2\sigma$), the RLT 61, and \mathfrak{P} 220 (in the original image resolution).

5.5.3 Potential cost image

In order to compute \mathfrak{P} we process the image so that minimal paths actually correspond to guide-wires. In this context the potential cost image is often derived from a Hessian based metric that already encodes information about curvilinear structures (see Section 5.3.1). In order to illustrate the power of the proposed technique we generate a very simple potential cost image based on scale selection. We remove background variations and select thin dark object with a dark top hat. The guide-wire radius being between 1 and 2 pixels, we use a square structuring element of 2 pixels in radius. In the resulting image the guide-wire is actually a path that is of low intensity. Summing the pixels along the path and retaining the minimal path, as proposed for the construction of the polygonal path image shall thus enable to select it. However, its tip is much more contrasted and tortuous than its body and tends to better attract path than the body does. In order to get rid of this unwanted behavior we simply threshold the image to isolate the tip, and set the corresponding pixels to a more moderate value.

5.5.4 Qualitative results

We illustrated the result of the voting technique ($\vartheta_{\tau_{\min}}(\mathfrak{P})$) on the last frame of each image sequence of our database in Fig. 5.35, 5.36, 5.37 and 5.38. We observe that the results are generally speaking very satisfactory: most of the guide-wire is enhanced, and few false alarm remain in the background. However three problems remain. The guide-wire tip attacks paths too much (creating false positives in its vicinity) and is too tortuous to be accurately fitted by our detector with the same setting as the guide-wire body. We shall detect it separately. Secondly, some linear structures are detected in the background that are not guide-wires. Since they satisfy all the properties that we selected for guide-wires, a higher level processing, based on other criteria shall handle them. For instance the presence of a tip is very characteristic of the guide-wire, as well a the motion that animate it. Finally, \mathfrak{P} is not able to handle some shapes properly, especially S shapes since they do not stay in the cones over long distances. In such cases the results are not as good as on simpler guide-wire shapes.




FIGURE 5.35: Application of the PSI to clinical examples.

5.5.5 Quantitative results

We used the ROC curve formalism described in Section 5.3.3.2 to quantify the results obtained on the database of clinical cases with four different methods: $\vartheta(\mathfrak{P})$, $\vartheta_{\tau_{\min}}(\mathfrak{P})$, and two classical detectors, a Rotated Line Template (RLT) [BVG09] and Hessian based techniques (Frangi's vesselness [FNVV98]). We have computed the partial AUC independently on each image of each sequence. We report in Fig. 5.39, for each sequence s the mean and standard deviation of the AUC denoted $AUC_{\mu}(s)$ and $AUC_{\sigma}(s)$. The




FIGURE 5.36: Application of the PSI to clinical examples.

height of the columns is the mean and the error bars represent mean plus/minus one standard deviation. We observe that when the whole image is considered for the computation of false positives, the performance of classical detectors decreases dramatically. This is due to the fact that they detect a large proportion of anatomical structures in the background that yield larger responses than the guide-wire. These structures are not impacting \mathfrak{P} because they are attenuated by the top hat in the potential cost image and because they cannot be linked into long smooth curvilinear structures. In order not to




FIGURE 5.37: Application of the PSI to clinical examples.

take these structures into account in our assessment we restrained the computation of false positives to a band around the guide-wire (see Section 5.3.3.2). We observe that on most of the sequences the line enhancement techniques derived from \mathfrak{P} are better than the classical ones. This general trend is illustrated on Fig. 5.40 that illustrate the mean and standard deviation of $AUC_{\mu}(s)$ over all the sequences, denoted AUC_{μ} and AUC_{σ} (see Eq. 5.36).




FIGURE 5.38: Application of the PSI to clinical examples.

$$AUC_{\mu} = \frac{1}{n} \sum_s AUC_{\mu}(s) \quad (5.35)$$

$$AUC_{\sigma} = \sqrt{\frac{1}{n} \sum_s (AUC_{\mu}(s) - AUC_{\mu})^2} \quad (5.36)$$




FIGURE 5.39: Illustration of the performance of the line enhancement techniques (measured by the partial AUC) on each sequence of the database. Top: FPR is computed on the whole image. Bottom: FPR is computed only in a band around the guide-wire.




FIGURE 5.40: Average performance of four line enhancement techniques over our database of clinical cases. Top: FPR is computed on the whole image. Bottom: FPR is computed only in a band around the guide-wire.

We can sort the techniques in two groups based on their performance. The first one, constituted of the line enhancement techniques derived from \mathfrak{P} performs better than the second one formed of classical techniques (RLT and Hessian). The difference between the two groups is greater than the differences inside each group. Local descriptors like RLT

and Hessian show variations of a few percent whereas \mathfrak{P} shows twice their performance. Moreover we note that, in accordance with other publications [ZABT04, BVG09, Tan10], the RLT performs better than Frangi’s vesselness. We have previously discussed this fact and proposed an explanation in Section 5.3.3.3. Finally the results show that path selection based on smoothness has a positive impact in this application. We can propose the following ranking of the techniques:

$$\vartheta_{\tau_{\min}}(\mathfrak{P}) > \vartheta(\mathfrak{P}) \gg \text{RLT} > \text{Vesselness} \quad (5.37)$$

In order to illustrate the performance of the different techniques we selected an ”optimal” threshold for which we display the result (see Section 5.3.3.2). We adopt the metrics of Missed Detection Rate (MDR) and False Detection Rate (FDR) presented in [HVP10a] to define this optimal threshold. When the threshold increases the MDR increases and the FDR decreases. Since we want both to be small we select the threshold minimizing MDR+FDR. For this experiment we restricted the computation of FDR to a band around the guide-wire. Fig. 5.41 displays the evolution of MDR, FDR and their sum as function of the threshold on three different sequences. We illustrate the thresholded images obtained with the four methods in Fig. 5.42, 5.43, 5.44 and 5.46. The outcomes are well aligned with the conclusions drove from the partial AUC: \mathfrak{P} based methods perform significantly better than the other ones. The smoothness constraints in $\vartheta_{\tau_{\min}}(\mathfrak{P})$ removes some anatomical false alarms that yield tortuous paths, and RLT is superior to Frangi’s vesselness.




FIGURE 5.41: MDR (blue), FDR (red) and MDR+FDR (green) as a function of the threshold in three different clinical cases. As expected, MDR increases, FDR decreases and their sum has a clear minimum. In this experiment we chose 500 thresholds regularly sampled between the max and the min of the image.

5.6 Conclusion and further work

We have presented in this chapter a reflexion on curvilinear structure enhancement in image processing that led us to propose an efficient new approach, the polygonal path image (\mathfrak{P}) that has several interesting properties:

- an efficient computational scheme,




FIGURE 5.42: Top: input image. Middle row: Frangi (left) and RLT (right). Bottom row: $\vartheta(\mathfrak{P})$ (left) and $\vartheta_{\tau_{\min}}(\mathfrak{P})$ (right). Observe that traditional methods (middle row) enhance a large amount of non relevant structures in the background. Moreover, they fail to enhance some parts of the guide-wire. These problems are not present with $\vartheta_{\tau_{\min}}(\mathfrak{P})$.

- the ability to control the smoothness and length of the structures to be analyzed,
- a unification of local, semi local, and global curvilinear structure analysis in a single framework,
- intuitive and simple parametrization.




FIGURE 5.43: Similar experiment as Fig. 5.42 on another clinical case. This one is the most challenging case yielding the poorest results due to the presence of the ribs that are very visible under this viewing angle and of sternal wires (first sequence in Fig. 5.39). Same observations apply.

This structure is a very rich descriptor of the curvilinear structures present in the image from which we derived several tools: line enhancement and direction fields computation techniques. Contrary to the direction fields that are built with state of the art methods that are non relevant outside of the curvilinear structures, ours indicate the direction of the closest such structure. We have demonstrated the relevance of our line enhancement technique for guide-wire segmentation both qualitatively and quantitatively in the ROC analysis formalism and compared it to state of the art techniques.




FIGURE 5.44: Similar experiment as Fig. 5.42 on another clinical case. Same observations apply.

We also believe that the reflexion we proposed on line enhancement comparison methodologies could be useful to other researchers.

Future work may include some natural extensions of the usage/construction of \mathfrak{P} that include:

- compute rank order costs instead of average along the paths.




FIGURE 5.45: Similar experiment as Fig. 5.42 on another clinical case. Same observations apply.

- extend the computation of \mathfrak{P} to vector potentials to better constrain local orientation,
- include a tortuosity cost in the cost of paths to build \mathfrak{P} ,
- iterate \mathfrak{P} ,
- extend the set of operators on the path structure,




FIGURE 5.46: Similar experiment as Fig. 5.42 on another clinical case. Same observations apply.

- select the optimal path length at each pixel.

Finally we believe that \mathfrak{P} is a generic tool for the analysis and processing of curvilinear structure that can find applications for the filtering and tracking of such structures for instance.

Chapter 6

Conclusion and perspectives

This PhD thesis falls into the context of curing coronary heart disease with minimally invasive techniques. These procedures involve a clinician inserting interventional tools inside the patient's vasculature, under the monitoring of Xray fluoroscopy. The purpose of the intervention is to restore the blood flow in a narrowed coronary artery by delivering an endo-prosthesis, called a stent. The objective of the work presented in this thesis was to design, study and validate, image processing techniques to improve the visualization of stents with the Xray imaging system. To this aim, we studied the automated detection, tracking and registration of two closely related interventional tools: the balloon markerballs and the guide-wire. Moreover, our work on guide-wire detection led us to propose a new generic tool for the processing of curvilinear structures: the *polygonal path image*.

In order to make the use of stent visualization enhancement techniques possible in daily clinical practice, our objective was twofold:

- yield significant improvements in image quality,
- reach a high level of automation and robustness.

In this conclusion, we first summarize the main achievements of our work that we regroup under the following topics: 2D and 3D stent visualization, image processing tools for interventional imaging, curvilinear structure processing and methodology. Some of these developments have been published and/or patented (refer to my publication list in App. A), some appear publicly for the first time here (mainly the content of Chapters 3, 4 and 5). We then recall the identified limitations of the proposed techniques and expose perspectives to overcome them.

6.1 Main results

6.1.1 2D and 3D stent visualization enhancement

In Chapter 1, we presented the interventional imaging system, the angioplasty procedures and the existing stent visualization techniques. Based on this context, we summarized the needs and requirements of the clinician regarding stent visualization. In Chapter 2, we pursued this analysis and translated the clinician needs into quantifiable evaluation metrics and algorithm design-guidelines. We proposed in Chapters 2, 3 and 4, a set of algorithms for 2D and 3D stent visualization. They involved the design of fully automated techniques, that include two original contribution in our field: non-linear registration and guide-wire subtraction. Regarding the 2D case, we presented a comprehensive study of the stent visualization problem covering thoroughly the technical and clinical aspects of this topic. We demonstrated the following results:

- Our method yields an improvement in image quality over the input images of more than 1 point on a 5 level scale (p value < 0.001), as assessed by two clinicians on nearly 200 clinical cases.
- Our non-linear registration is superior to the state of the art similarity transform (improvement is visible in 27% of the cases). It has been evaluated by the author of this thesis on 140+ cases.
- The overall approach that we proposed, has been compared blindly to another commercially available software by three clinicians. Our technique has been judged to yield superior results (p value < 0.001 - A noticeable improvement was observed in 74% of the cases).
- The guide-wire that impairs the visibility of stent details can be removed from the image with image processing techniques without degrading the stent information (illustrated on 10+ clinical cases).

Regarding 3D stent visualization, we demonstrated on a database of 22 clinical cases:

- A fully automated algorithm to perform 3D stent reconstruction that yields no visual difference with the reference, non-automated, reconstruction.
- The extension to the 3D case of the tools that we developed for 2D stent visualization (non-linear registration and guide-wire removal), yielding promising results.

6.1.2 Image processing tools for interventional imaging

The stent visualization techniques that we developed and exposed in this thesis involved the design and validation of the following set of algorithms:

- an algorithm to detect and track the angioplasty balloon markerballs in short (≈ 30 frames), static, image sequences. This fully algorithm has been validated on 7000+ images. Its success rate is superior to 90% and its accuracy on the position of the markerballs is sub-pixel.
- a similar algorithm for the case of long rotational image sequences (> 100 images). Its evaluation on a database of ≈ 3000 images demonstrated a success rate of 98%. These results are the best reported so far in this field.
- a technique to segment the guide-wire between two given endpoints. In typical clinical conditions, its precision is inferior to the radius of the guidewire in 97% of the cases.
- a non-linear registration technique, tailored for the registration of stents and vessels.
- a guide-wire removal technique that does not degrade the content of the image that is overlapping with the guide-wire.

6.1.3 Curvilinear structure processing

In Chapter 5 we were interested in the segmentation of guide-wires without any user interaction or *a priori* on the presence of markerballs (*i.e.* without given endpoints for instance). In this context we developed a tool to process curvilinear structures, denoted the polygonal path image \mathfrak{P} , that has several interesting properties :

- a unification of local, semi local, and global curvilinear structure analysis in a single framework,
- the ability to control the smoothness and length of the structures to analyze,
- an intuitive and simple parametrization,
- an efficient computational scheme.

We derived from \mathfrak{P} several curvilinear structure processing techniques. We demonstrated the superiority of this tool over state of the art line enhancement techniques on a database of 100+ clinical sequences: Its performance is twice better than the state of the art, as measured by the partial area under the ROC curve.

6.1.4 Methodology

Each of the quantitative validation that we summarized in the previous sections of this conclusion, required to design an appropriate measurement methodology. We proposed validation methodologies to quantify the performance of algorithms for the following tasks:

- automatic markerball segmentation,
- automatic guide-wire segmentation,
- relative and absolute visual quality assessment of stent images.

6.2 Limitations and technical perspectives

6.2.1 Limitations

6.2.1.1 Image quality of enhanced stent images

We noted that the stent visualization techniques presented in this thesis, despite their good performance versus state of the art ones, do not always enable the observation of all the desired stent details. We believe that one of the main contributors to this image quality limitation is the relative motions of the landmark versus the stent (that we just call “relative motion” in this paragraph). Rigidly linking the landmarks to the stent, in order to annihilate this motion, can be achieved by inflating the balloon. However, it is unlikely that clinicians would perform this action in daily practice, for the sake of visualizing the stent (it results in blocking the blood flow of the artery). Therefore, the development of image processing techniques that would compensate the relative motion is a potential field of further research.

A second potential solution to improve image quality in stent visualization applications, is to select a subset of the input image to take into account. This strategy allows to reject the frames where the relative motion is the largest.

6.2.1.2 Applicability of stent visualization techniques

All the stent visualization techniques that we have presented (Chapters 2, 3, and 4) rely on the presence of the guide-wire and of the markerballs of the angioplasty balloon. Although the presence of the guide-wire is natural since it is the first device to enter the artery, and the last one to leave it, the presence of the markerballs can be regarded as a constraint for the clinician. In several cases, he faces the desire to visualize a stent while there is no markerball in its close vicinity. The techniques that we propose do not allow to answer this need. A stent visualization techniques that would not require the presence of the markerball would better integrate into the clinical workflow and potentially save dose (since all the stents present in the artery could be enhanced with a single Xray sequence).

6.2.1.3 Co-visualization of stents and vessels

When we listed the needs of the clinicians regarding stent visualization in Chapter 1, we expressed the interest of visualizing the stent with respect to its environment: the vessel lumen, the potential bifurcations and the calcifications. The techniques that we presented focus on the stent and did not bring any information regarding the vessel itself. Techniques that enable the joint visualization of stents and vessels have recently appeared on the market [SF09, Inc11] and will probably be the subject of further research.

6.2.2 Perspectives

6.2.2.1 Layer separation

We presented a three-layer-model (background, guide-wire and stent) for the stent images, and proposed a method to extract the stent layer. The key point of the method is that the guide-wire can be accurately modeled by a parametric model. We believe that this formulation can potentially be extended to separate other structures in the images, provided that one can come up with a relevant parametric model. It could be used, for instance, to remove the markerballs or some simple anatomical features.

6.2.2.2 Curvilinear structure processing

The tool that we have presented for the processing of curvilinear structures has been used to derive enhancement and direction estimation algorithms. Our future work will include full guide-wire segmentation and the application of \mathfrak{P} to other structures in biomedical imaging. Moreover the theoretical foundation in which it fits needs to be further developed.

6.2.2.3 Interventional tool detection and tracking

As previously mentioned, the developments of Chapters 2 and 3 are the basis of the StentViz feature commercially available on the Innova systems. Other similar products exist on the market, that we regrouped under the generic name Digital Stent Enhancement (DSE). To the best of our knowledge DSE is the first family of industrial products based on the automatic detection and tracking of interventional tools. As mentioned in Section 1.4.1, the research in this area has been intensifying over the past years. Automatic detection and tracking of interventional tool is being studied for a large number of specific interventional tools (lasso catheters, ablation catheters, guide-wire, IVUS probes etc). As a result, we expect new clinical applications and products to appear in a near future. Finally, the real-time extension of the concept of DSE seems to be a logical next step towards improved guidance during PCI.

6.2.2.4 Clinical impact and applications

The DSE technique that we proposed, integrated into the Innova series (GE Healthcare), is installed on a large number of clinical sites. Beyond the technical benefits that we have demonstrated (workflow, image quality...), we hope that further long-term studies will ultimately demonstrate that some the tools we develop can yield a significant improvement in either:

- cost: reduce the cost of interventions.
- quality: improve patient outcome.
- access: enable to treat patients that could not benefit from these interventions.

Appendix A

Publication list

Some of the material described in this manuscript have been the subject of publications, patents and master of science thesis.

A.1 Journal

- V. Bismuth, R. Vaillant, F. Funck, N. Guillard, and L. Najman. A comprehensive study of stent visualization enhancement in x-ray images by image processing means. *Medical Image Analysis*, 2011

A.2 Conferences

- V. Bismuth, R. Vaillant, H. Talbot, L. Najman. Curvilinear structure enhancement with the polygonal path image - Application to guide-wire segmentation in Xray fluoroscopy. *MICCAI 2012*.
- F. Funck, T. Lefevre, V. Bismuth, and R. Vaillant. Guide-wire subtraction in enhanced stent images. *EuroIntervention*, Supplement H, 6, May 2010
- F. Funck, N. Guillard, V. Bismuth, and R. Vaillant. Quantitative evaluation of an elastic stent enhancement technique in x-ray image sequences. *EuroIntervention*, Supplement E, 5, May 2009
- V. Bismuth, L. Vancamberg, and S. Gorges. A comparison of line enhancement techniques: applications to guide-wire detection and respiratory motion tracking. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7259, feb 2009.
- V. Bismuth and R. Vaillant. A device enhancing and denoising algorithm for x-ray cardiac fluoroscopy. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pages 1-4. IEEE, 2008.

- V. Bismuth and R. Vaillant. Elastic registration for stent enhancement in x-ray image sequences. In International Conference on Image Processing (ICIP), pages 2400-2403, 2008.

A.3 Patents

We restrict the patent section to the ones directly related to the matter of this PhD.

- V. Bismuth and R. Vaillant. Procédé de traitement d'images radiologiques - french application number 1000078096, 2010
- R. Vaillant and V. Bismuth. Procédé de traitement d'images radiologiques - french application number 1000076021, 2010
- R. Vaillant, J. Liénard, and V. Bismuth. System and method to enhance visualization of an object in a vascular vessel - us patent 7734328, 2006.

A.4 Supervised master's thesis

I supervised the following Master of Science students for their final dissertations, that have close relations with the topics presented in this thesis:

- A. Chie. Reconstruction 3D de stents coronaires. Telecom ParisTech, 2008.
- P. Delafon. Visualisation de stents. Ecole Centrale Lyon, 2006.
- Y. Xi. Visualisation de stents. Telecom ParisTech, 2005.

I moreover supervised the following ones, that deal with the same problematic but whose technical developments have not been used directly in this thesis:

- R. Prevost. Covisualisation de stent et vaisseaux en imagerie cardiaque interventionnelle. ENSTA Paris, 2010.
- S. Blusseau. Covisualisation de stent et vaisseaux en imagerie cardiaque interventionnelle. Telecom ParisTech, 2010.
- Y. Levy. Réhaussement d'outils chirurgicaux en imagerie interventionnelle cardiaque. Telecom SudParis, 2009.
- J. Dura. Réhaussement d'outils chirurgicaux en imagerie interventionnelle cardiaque. Telecom SudParis, 2009.
- T. Courtat. Segmentation de séquences d'images en interventions cardiaques. Supelec, 2008.
- M. Gohniem. Suivi doutils interventionnels sur séquences dimages. ENSI Caen, 2007.

Appendix B

Resume

Work Experience

2010 - current General Electric Healthcare

Senior research engineer

In charge of developing new applications on the Innova product line (Image processing algorithm development, integration into software development cycles, publications, patents, clinical application expertise, clinical study management, coaching . . .)

2004 - 2010 General Electric Healthcare

Research engineer

2003 - 2004 General Electric Healthcare

Image quality engineer

2002 Thales optronic

Master's thesis internship : Image processing for automatic target detection and recognition.

Education

Dec 2008 - current Université Paris Est

PhD student in Computer Science under the supervision of Laurent Najman and Régis Vaillant

2001 - 2002 ENS Cachan

Master's degree : Mathematics, Vision, Learning - with high honors (mention très bien)

1999 - 2002 Ecole Centrale, Paris

MSc - Major in applied mathematics

1999 - 2000 Université Pierre et Marie Curie, Paris

BSc in applied mathematics - with high honors (mention très bien).

1998 - 1999 Lycée Charlemagnes, Paris

Post-secondary preparatory school / classes preparing for entrance examinations to the French Grandes Ecoles.

1997

1997: Baccalauréat S - Equivalent to English A-levels in pure maths and physics, Awarded with merit.

Teaching

2009 - current ESIEE Paris

Teaching Linear methods in image processing to Master students

Skills

Software languages

C/C++, Matlab, shell, Latex

Languages

English and Spanish: fluent

German, Portuguese and Hebrew: basic

Misc

6 sigma green belt certification (since 2003 @ GE)

> 15 patent applications

> 10 times Master thesis tutor

2008 - current @ Ecole Centrale Paris: Brazilian Jiu Jitsu instructor

Bibliography

- [ABL06] V. Auvray, P. Bouthemy, and J. Lienard. Motion estimation in x-ray image sequences with bi-distributed transparency. In Image Processing, 2006 IEEE International Conference on, pages 1057–1060. IEEE, 2006. [32](#)
- [ABL09] V. Auvray, P. Bouthemy, and J. Liénard. Jointmotion estimation and layer segmentation in transparent image sequences: application to noise reduction in X-ray image sequences. EURASIP Journal on Advances in Signal Processing, 2009:1–21, 2009. [84](#), [86](#)
- [AGZ⁺08] S. Atasoy, M. Groher, D. Zikic, B. Glocker, T. Waggershauer, M. Pfister, and N. Navab. Real-time respiratory motion tracking: roadmap correction for hepatic artery catheterizations. In Proceedings of SPIE, volume 6918, page 691815, 2008. [32](#), [135](#)
- [ALB05] V. Auvray, J. Liénard, and P. Bouthemy. Multiresolution parametric estimation of transparent motions and denoising of fluoroscopic images. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005, pages 352–360, 2005. [32](#)
- [Ans48] Francis John Anscombe. The transformation of Poisson, Binomial and Negative-Binomial data. Biometrika, 35(3-4):246–254, 1948. [36](#)
- [Ant08] E.M. Antman. Time is muscle:: Translation into practice. Journal of the American College of Cardiology, 52(15):1216–1221, 2008. [4](#)
- [AR07] F.J. Ayres and R.M. Rangayyan. Design and performance analysis of oriented feature detectors. Journal of Electronic Imaging, 16:023007, 2007. [142](#), [143](#), [145](#), [146](#)
- [ASPV⁺07] Fernando Alfonso, Alfonso Suarez, Maria J. Perez-Vizcayno, Raul Moreno, Javier Escaned, Camino Banuelos, Pilar Jimenez, Esther Bernardo, Dominick J. Angiolillo, Rosana Hernandez, and Carlos

- Macaya. Intravascular Ultrasound Findings During Episodes of Drug-Eluting Stent Thrombosis. J Am Coll Cardiol, page j.jacc.2007.08.015, 2007. [30](#)
- [ASS99] Til Aach, Ulrich Schiebel, and Gerhard Spekowius. Digital image acquisition and processing in medical x-ray imaging. Journal of Electronic Imaging, 8(Special Section on Biomedical Image Representation):7–22, 1999. [36](#)
- [B⁺67] H. Blum et al. A transformation for extracting new descriptors of shape. Models for the perception of speech and visual form, 19(5):362–380, 1967. [53](#)
- [BAG⁺07] Adrian Barbu, Vassilis Athitsos, Bogdan Georgescu, Stefan Böhm, Peter Durlak, and Dorin Comaniciu. Hierarchical learning of curves application to guidewire localization in fluoroscopy. In CVPR, 2007. [16](#), [50](#), [134](#), [135](#), [136](#), [137](#), [142](#), [143](#), [144](#)
- [BBHP92] J.R. Bergen, P.J. Burt, R. Hingorani, and S. Peleg. A three-frame algorithm for estimating two-component image motion. IEEE Transactions on Pattern Analysis and Machine Intelligence, pages 886–896, 1992. [84](#), [87](#)
- [BC09] F. Benmansour and L.D. Cohen. Fast object segmentation by growing minimal paths from a single point on 2d or 3d images. Journal of Mathematical Imaging and Vision, 33(2):209–221, 2009. [xxvi](#), [171](#)
- [BCD⁺07] A.J. Boyle, M.Y. Chan, J. Dib, N.K. Kapur, S. Kraft, R. Vaillant, J.S. Whiting, and J.O.N.R. Resar. Assessment of a novel angiographic image stabilization system for percutaneous coronary intervention. Journal of interventional cardiology, 20(2):153–157, 2007. [32](#)
- [BD99] A. Brandt and J. Dym. Fast calculation of multiple line integrals. SIAM J. Sci. Comput., 20:1417–1429, February 1999. [156](#)
- [BDJ⁺04] OF Bertrand, R DeLarochelli re, M Joyal, R Bonan, R Mongrain, and JC Tardif. Incidence of stent under-deployment as a cause of in-stent restenosis in long stents. Int J Cardiovasc Imaging, 20(4):279–84, 2004. [30](#)
- [BDN09] M. Baust, S. Demirci, and N. Navab. Stent graft removal for improving 2d–3d registration. In Biomedical Imaging: From Nano to Macro, 2009. ISBI’09. IEEE International Symposium on, pages 1203–1206. IEEE, 2009. [82](#)

- [BG09] Vincent Bismuth and Sebastien Gorges. Shape from radio-opacity for interventional devices 3d, 2009. French Patent number 0955984. [89](#), [96](#)
- [BK95] J.C. Brailean and A.K. Katsaggelos. Simultaneous recursive displacement estimation and restoration of noisy-blurred image sequences. Image Processing, IEEE Transactions on, 4(9):1236–1251, 1995. [32](#)
- [BKE⁺95] J.C. Brailean, R.P. Kleihorst, S. Efstratiadis, A.K. Katsaggelos, and R.L. Lagendijk. Noise reduction filters for dynamic image sequences: a review. Proceedings of the IEEE, 83(9):1272–1292, 1995. [32](#)
- [BLA02] C Brasselet, A Lafont, and F Addad. Qu'est-ce qu'une sténose coronaire significative ? La Lettre du cardiologue, 352:32–36, 2002. [4](#)
- [BLHS09] A. Brost, R. Liao, J. Hornegger, and N. Strobel. 3-d respiratory motion compensation during ep procedures by image-based 3-d lasso catheter model generation and tracking. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009, pages 394–401, 2009. [16](#), [32](#), [135](#), [136](#), [137](#)
- [BLSH10] A. Brost, R. Liao, N. Strobel, and J. Hornegger. Respiratory motion compensation by model-based catheter tracking during ep procedures. Medical image analysis, 14(5):695–706, 2010. [32](#), [135](#)
- [BMVA06] Christophe Blondel, Gregoire Malandain, Regis Vaillant, and Nicholas Ayache. Reconstruction of coronary arteries from a single rotational x-ray projection sequence. IEEE transactions on medical imaging, pages 653–663, 2006. [123](#)
- [Bra97] A.P. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. Pattern Recognition, 30(7):1145–1159, 1997. [139](#), [142](#)
- [BSCB00] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In Proceedings of the 27th annual conference on Computer graphics and interactive techniques, pages 417–424. ACM Press/Addison-Wesley Publishing Co., 2000. [82](#)
- [Bus02] J.T. Bushberg. The essential physics of medical imaging. Williams & Wilkins, 2002. [83](#)
- [BV08a] V. Bismuth and R. Vaillant. A device enhancing and denoising algorithm for x-ray cardiac fluoroscopy. In Pattern Recognition, 2008.

- ICPR 2008. 19th International Conference on, pages 1–4. IEEE, 2008. [16](#)
- [BV08b] Vincent Bismuth and Régis Vaillant. Elastic registration for stent enhancement in x-ray image sequences. In International Conference on Image Processing (ICIP), pages 2400–2403, 2008. [18](#), [20](#), [29](#), [31](#), [53](#), [64](#)
- [BV10] Vincent Bismuth and Régis Vaillant. Procédé de traitement d’images radiologiques - french application number 1000078096, 2010. [18](#), [21](#), [82](#)
- [BVF⁺11] V. Bismuth, R. Vaillant, F. Funck, N. Guillard, and L. Najman. A comprehensive study of stent visualization enhancement in x-ray images by image processing means. Medical Image Analysis, 2011. [18](#), [20](#), [24](#), [29](#)
- [BVG09] V. Bismuth, L. Vancanberg, and S. Gorges. A comparison of line enhancement techniques: applications to guide-wire detection and respiratory motion tracking. In Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, volume 7259, feb 2009. [16](#), [18](#), [32](#), [50](#), [134](#), [135](#), [136](#), [142](#), [143](#), [145](#), [146](#), [147](#), [174](#), [179](#)
- [BVN03] S.A.M. Baert, M.A. Viergever, and W.J. Niessen. Guide-wire tracking during endovascular interventions. Medical Imaging, IEEE Transactions on, 22(8):965–972, 2003. [134](#), [135](#), [136](#), [137](#)
- [BWL⁺10] Alexander Brost, Andreas Wimmer, Rui Liao, Joachim Hornegger, and Norbert Strobel. Catheter tracking: Filter-based vs. learning-based. In Michael Goesele, Stefan Roth, Arjan Kuijper, Bernt Schiele, and Konrad Schindler, editors, Pattern Recognition, volume 6376 of Lecture Notes in Computer Science, pages 293–302. Springer Berlin / Heidelberg, 2010. [16](#), [135](#)
- [CAC⁺09] J Córdova, G Aleong, H Colmenarez, A Cruz, E Canales, P Jimenez-Quevedo, R Hernández, F Alfonso, C Macaya, C Banuelos, W DenHartog, and J Escaned. Digital enhancement of stent images in primary and secondary percutaneous coronary revascularisation. EuroIntervention : journal of EuroPCR, Suppl D:D101-6, may 2009. [31](#), [73](#), [75](#)
- [CAEB04] A. Condurache, T. Aach, K. Eck, and J. Bredno. Fast detection and processing of arbitrary contrast agent injections in coronary angiography and fluoroscopy. Bildverarbeitung fuer die Medizin 2004 (Algorithmen, Systeme, Anwendungen), pages 5–9, 2004. [88](#)

- [Car07] Mark J. Carlotto. Enhancement of low-contrast curvilinear features in imagery. IEEE Transactions on Image Processing, 16(1):221–228, 2007. 158
- [CAW00] R. A. Close, C. K. Abbey, and J. S. Whiting. Improved image guidance of coronary stent deployment. In S. K. Mun, editor, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, volume 3976, pages 301–304, April 2000. 16, 31, 32, 65, 73
- [CAW02] R. A. Close, C. K. Abbey, and J. S. Whiting. Improved Localization of Coronary Stents Using Layer Decomposition . In Computer Aided Surgery, Issue 2, volume 7 of Computer Aided Surgery, pages 84 – 89, 2002. 17, 73, 84, 85
- [Chi06] Grahams Child. Endo dysfunction athero.png - this file is licensed under the creative commons attribution-share alike 3.0 unported license. http://en.wikipedia.org/wiki/File:Endo_dysfunction_Athero.PNG, 2006. xv, 5
- [Cho10] Cheol Ung Choi. Tct-375: Impact of stentboost enhancement guided percutaneous coronary intervention on midterm angiographic and clinical outcomes. J Am Coll Cardiol, 56(13 MeetingAbstracts):B87–a–, 2010. 75
- [Chu92] C.K. Chui. An introduction to wavelets, volume 1. Academic Pr, 1992. 146
- [CJHJ05] R. Carrillo-Jimenez, S.L. Houser, and F.A. Jaffer. Culprit lesion atherothrombectomy during acute myocardial infarction: Extraction of an acute coronary plaque rupture. Circulation, 112(16):e267, 2005. xvii, 27
- [CK97] L.D. Cohen and R. Kimmel. Global minimum for active contour models: A minimal path approach. International Journal of Computer Vision, 24(1):57–78, 1997. 149, 166
- [CKS93] C.L. Chan, A.K. Katsaggelos, and A.V. Sahakian. Image sequence filtering in quantum-limited noise with applications to low-dose fluoroscopy. Medical Imaging, IEEE Transactions on, 12(3):610–621, 1993. 32
- [CM09] A. Condurache and A. Mertins. A point-event detection algorithm for the analysis of contrast bolus in fluoroscopic images of the coronary

- arteries. In 17th European Signal Processing Conference (EUSIPCO 2009), pages 2337–2341, 2009. 88
- [CMW97] Robert A. Close, Craig A. Morioka, and James S. Whiting. Fluoroscopic 3d reconstruction of coronary stents. volume 3032, pages 445–450. SPIE, 1997. 102, 130
- [CPT03] A. Criminisi, P. Perez, and K. Toyama. Object removal by exemplar-based inpainting. In Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on, volume 2, pages II–721. IEEE, 2003. 82
- [CSE06] I.K. Cihan, H.G. Senel, and A.S. Havelsan Ehsim. An application of topological median on detection and clustering of microcalcification in digital mammograms. IEEE International Conference on Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings, 2:II–II, May 2006. 48
- [CSMS05] DSG Conway, WHT Smith, J Moore, and U Sivananthan. Measurement of coronary stent expansion using stentboost image enhancement software: a comparison with intravascular ultrasound. British Heart Journal, 91 (supplement I):A39–A40, 2005. 31, 75
- [CWA03] Robert A. Close, James S. Whiting, and Craig K. Abbey. Image guidance for coronary stent deployments, 2003. US Patent 6532380. 17, 73
- [DAJW10] A. Dixon, R. Areste, K.N. Jabri, and V. Walimbe. Novel registration-based image enhancement for x-ray fluoroscopy. In Proceedings of SPIE, volume 7623, page 76233D, 2010. 32
- [DC09] Diagnostic and Interventional Cardiology. Boston scientific launches taxus liberte des in japan. <http://www.rsnafastpass.net/node/31235/3>, 2009. xv, 10
- [DC10a] Diagnostic and Interventional Cardiology. Brakes on cta and mra. <http://www.rsnafastpass.net/node/36125/>, 2010. xv, 14
- [DC10b] Diagnostic and Interventional Cardiology. Drug-eluting stents help patients avoid amputation. <http://www.dicardiology.com/article/drug-eluting-stents-help-patients-avoid-amputation>, 2010. xv, 10

- [DCOY03] I. Drori, D. Cohen-Or, and H. Yeshurun. Fragment-based image completion. In ACM SIGGRAPH 2003 Papers, pages 303–312. ACM, 2003. 82
- [DH72] R.O. Duda and P.E. Hart. Use of the hough transformation to detect lines and curves in pictures. Communications of the ACM, 15(1):11–15, 1972. 172
- [DHJ⁺01] David L. Donoho, Xiaoming Huo, Ian Jermyn, Peter Jones, Gilad Lerman, Ofer Levi, and Frank Natterer. Beamlets and multiscale image analysis. Multiscale and Multiresolution Methods: Theory and Applications, 2001. 156
- [DKT05] VP Dinesh Kumar and T. Thomas. Automatic estimation of orientation and position of spine in digitized x-rays using mathematical morphology. Journal of Digital Imaging, 18(3):234–241, 2005. 88
- [DS84] E. Dubois and S. Sabri. Noise reduction in image sequences using motion-compensated temporal filtering. Communications, IEEE Transactions on, 32(7):826–831, 1984. 32
- [Ele11] ”General Electric”. Innova (ge website). <https://www2.gehealthcare.com/portal/site/usen/menuitem.d9d1e5260a507013d6354a1074c84130/?vgnnextoid=ca8c9daa2d930210VgnVCM10000024dd1403RCRD&vgnnextfmt=default>, 2011. 9
- [ENS⁺10] Andrejs Erglis, Inga Narbutė, Dace Sondore, Alona Grave, and Sanda Jegere. Tools and techniques: Coronary guidewires. EuroInterventions, (6):1–8, 2010. xx, 89
- [esa11] esaote. Esaote - suitestensa for cardiology. http://www.esaote.com/modules/core/page.asp?p=SUITESTENSA_CARDIO&t=CATHLAB_M, 2011. xv, 6
- [FAoTMLVG91] W.T. Freeman, E.H. Adelson, Massachusetts Institute of Technology. Media Laboratory. Vision, and Modeling Group. The design and use of steerable filters. IEEE Transactions on Pattern analysis and machine intelligence, 13(9):891–906, 1991. 136, 146
- [Faw06] T. Fawcett. An introduction to roc analysis. Pattern recognition letters, 27(8):861–874, 2006. 139, 142, 145

- [FCM⁺05] Kenichi Fujii, Stephane G. Carlier, Gary S. Mintz, Yi-ming Yang, Issam Moussa, Giora Weisz, George Dangas, Roxana Mehran, Alexandra J. Lansky, Edward M. Kreps, Michael Collins, Gregg W. Stone, Jeffrey W. Moses, and Martin B. Leon. Stent underexpansion and residual reference segment stenosis are related to stent thrombosis after sirolimus-eluting stent implantation: An intravascular ultrasound study. *J Am Coll Cardiol*, 45(7):995–998, 2005. [30](#)
- [FG87] W. Forstner and E. Gulch. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proc. ISPRS Intercommission Conference on Fast Processing of Photogrammetric Data*, pages 281–305, 1987. [146](#)
- [FGBV09] F Funck, N Guillard, V Bismuth, and R Vaillant. Quantitative evaluation of an elastic stent enhancement technique in x-ray image sequences. *EuroIntervention, Supplement E*, 5, May 2009. [xvi](#), [18](#), [19](#), [20](#), [29](#), [31](#), [59](#), [75](#)
- [FLBV10] F Funck, Thierry Lefevre, V Bismuth, and R Vaillant. Guide-wire subtraction in enhanced stent images. *EuroIntervention, Supplement H*, 6, May 2010. [18](#), [21](#), [79](#), [82](#), [93](#)
- [FNL08] Raoul Florent, Lucile Nosjean, Pierre Lelong, and Peter Maria Johannes Rongen. Medical viewing system and method for enhancing structures in noisy images, 2008. US Patent 7415169. [16](#), [17](#), [31](#), [32](#), [50](#), [65](#), [73](#)
- [FNVV98] Alejandro F. Frangi, Wiro J. Niessen, Koen L. Vincken, and Max A. Viergever. Multiscale vessel enhancement filtering. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI) 1998*, volume 1496 of *Lecture Notes in Computer Science*, pages 130–. 1998. [136](#), [146](#), [147](#), [172](#), [174](#)
- [FRvAtHR06] E. Franken, P. Rongen, M. van Almsick, and B. ter Haar Romeny. Detection of electrophysiology catheters in noisy fluoroscopy images. In *Proc. MICCA, LNCS*, volume 4191, pages 25–32, 2006. [134](#), [135](#), [136](#), [137](#)
- [FYL⁺03] W.F. Fearon, A.C. Yeung, D.P. Lee, P.G. Yock, and P.A. Heidenreich. Cost-effectiveness of measuring fractional flow reserve to guide coronary interventions. *American Heart Journal*, 145(5):882–887, 2003. [8](#)

- [GBB⁺98] C.J. George, D.S. Baim, J.A. Brinker, D.L. Fischman, S. Goldberg, R. Holubkov, E.D. Kennard, L. Veltri, and K.M. Detre. One-year follow-up of the stent restenosis (stress i) study 1. The American journal of cardiology, 81(7):860–865, 1998. 8
- [GBK⁺09] M. Groher, F. Bender, A. Khamene, W. Wein, T. H. Heibel, and N. Navab. 3D Guide Wire Navigation from Single Plane Fluoroscopic Images in Abdominal Catheterizations, pages 356–+. Springer, 2009. 16, 135
- [GDLA04] C.M. Gibson, J.A. De Lemos, and E.M. Antman. Time is muscle in primary pci: the strength of the evidence grows. European heart journal, 25(12):1001, 2004. 4
- [GLA⁺04] P. Greenland, L. LaBree, S.P. Azen, T.M. Doherty, and R.C. Detrano. Coronary artery calcium score combined with framingham score for risk prediction in asymptomatic individuals. JAMA: the journal of the American Medical Association, 291(2):210, 2004. 13
- [GM91] Z. Gigus and J. Malik. Detecting curvilinear structure in images. UC Berkeley CSD Technical Report, 91:619, 1991. 136, 137
- [Gri91] Michel Grimaud. La géodésie numérique en morphologie mathématique. Application à la détection automatique des microcalcifications - 2eme partie. PhD Thesis. PhD thesis, "Ecole des Mines de Paris", 1991. 48
- [GS10a] S. Garg and P.W. Serruys. Coronary stents:: Looking forward. Journal of the American College of Cardiology, 56(10):S43–S78, 2010. 27
- [GS10b] Scot Garg and Patrick W. Serruys. Coronary stents: Current status. J Am Coll Cardiol, 56(10 Suppl S):S1–42, 2010. 8, 9, 12, 15, 30
- [HGG⁺09] T.H. Heibel, B. Glocker, M. Groher, N. Paragios, N. Komodakis, and N. Navab. Discrete tracking of parametrized curves. Computer Vision and Pattern Recognition, IEEE Computer Society Conference on, 0:1754–1761, 2009. 16, 50, 135, 136, 137
- [HKR93] D.P. Huttenlocher, G.A. Klanderman, and WA Rucklidge. Comparing images using the hausdorff distance. IEEE Transactions on pattern analysis and machine intelligence, pages 850–863, 1993. 162
- [Hos07] European Hospital. Copyright 1997-2009 by european hospital verlags gmbh. <http://www.european-hospital.com/en/article/>

- [2455-E-SIRIUS_Trial_Safety_and_efficacy_of_drug-eluting_stents.html](#), 2007. xv, 10
- [HTD05] S. Hadjidemetriou, D. Toomre, and J.S. Duncan. Segmentation and 3d reconstruction of microtubules in total internal reflection fluorescence microscopy (tirfm). Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005, pages 761–769, 2005. 134, 135, 136, 137
- [HVP10a] N. Honnorat, R. Vaillant, and N. Paragios. Guide-wire extraction through perceptual organization of local segments in fluoroscopic images. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010, pages 440–448, 2010. 16, 50, 135, 136, 137, 142, 143, 144, 179
- [HVP10b] N. Honnorat, R. Vaillant, and N. Paragios. Robust guidewire segmentation through boosting, clustering and linear programming. In Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on, pages 924–927. IEEE, 2010. 16, 50, 135, 136, 142, 143, 144, 146
- [Inc11] ”Paieon Inc.”. Stentoptimize. <http://www.paieon.com/Products.asp?Par=9.27&id=75>, 2011. 189
- [Ind09] Indon. Roc space - this file is licensed under the creative commons attribution-share alike 3.0 unported license. http://en.wikipedia.org/wiki/File:ROC_space-2.png, 2009. xxiv, 141
- [IRP94] M. Irani, B. Rousso, and S. Peleg. Computing occluding and transparent motions. International Journal of Computer Vision, 12(1):5–16, 1994. 84, 87
- [JM02] X. Jiang and D. Mojon. Supervised evaluation methodology for curvilinear structure detection algorithms. Pattern Recognition, 1:10103, 2002. 144
- [JU04] M. Jacob and M. Unser. Design of steerable filters for feature detection using canny-like criteria. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 26(8):1007–1019, 2004. 136, 146
- [KBR⁺09] AP King, R. Boubertakh, KS Rhode, YL Ma, P. Chinchapatnam, G. Gao, T. Tangcharoen, M. Ginks, M. Cooklin, JS Gill, et al. A subject-specific technique for respiratory motion correction in image-guided cardiac catheterisation procedures. Medical image analysis, 13(3):419–431, 2009. 32

- [KJY⁺05] Vakharia Kalpesh, Mishell Jacob, Yeghiazarians Yerem, Ports Thomas, and Michaels Andrew. Tct 2005 abstract: Determination of adequate stent expansion: A comparison of quantitative coronary angiography, intravascular ultrasound, and novel stentboost x-ray imaging. J Am Coll Cardiol, 96:1H–213H, 2005. [75](#)
- [KL05] J.P. Kaufhold and D.A. Langan. Methods and appartus for visualizing low contrast moveable objects, 2005. US Patent US20050074158A1. [16](#), [31](#), [38](#), [74](#)
- [Koo05] J.J. Koolen. Stentboost image enhancement: first clinical experience. MedicaMundi, aug 2005. [31](#), [37](#), [38](#), [65](#), [73](#), [74](#), [75](#)
- [KS05] D. Kunz and B. Schweiger. Line detection in strongly noise-corrupted images. Bildverarbeitung für die Medizin 2005, pages 50–54, 2005. [136](#)
- [KSW01a] K.S. Kump, P.B. Sachs, and D.L. Wilson. Digital subtraction peripheral angiography using image stacking: Initial clinical results. Medical Physics, 28:1482, 2001. [38](#), [39](#)
- [KSW01b] K.S. Kump, G.M. Saidel, and D.L. Wilson. Comparison of algorithms for combining x-ray angiography images. Medical Imaging, IEEE Transactions on, 20(8):742–750, 2001. [38](#), [39](#)
- [KtB96] N. Karssemeijer and G.M. te Brake. Detection of stellate distortions in mammograms. Medical Imaging, IEEE Transactions on, 15(5):611–619, 1996. [146](#)
- [LABFL09] D. Lesage, E.D. Angelini, I. Bloch, and G. Funka-Lea. A review of 3d vessel lumen segmentation techniques: Models, features and extraction schemes. Medical Image Analysis, 13(6):819–845, 2009. [134](#), [135](#)
- [LAM⁺81] LA Lehmann, RE Alvarez, A. Macovski, WR Brody, NJ Pelc, SJ Riederer, and AL Hall. Generalized image combinations in dual kvp digital radiography. Medical physics, 8:659, 1981. [87](#)
- [LBL⁺09] S. Lessard, P. Bigras, C. Lau, D. Roy, G. Soulez, and J.A. de Guise. Automatically driven vector for guidewire segmentation in 2d and biplane fluoroscopy. World Academy of Science, Engineering and Technology, 54:446–450, 2009. [135](#), [136](#), [137](#)
- [Lia10] R. Liao. Model-based lasso catheter tracking in monoplane fluoroscopy for 3d breathing motion compensation during ep procedures. In Proceedings of SPIE, volume 7625, page 76250Y, 2010. [32](#), [135](#)

- [Lin96] T. Lindeberg. Edge detection and ridge detection with automatic scale selection. In *cvpr*, page 465. Published by the IEEE Computer Society, 1996. 146
- [Lin98] Tony Lindeberg. Feature detection with automatic scale selection. In *International Journal of Computer Vision*, volume 30 of *Lecture Notes in Computer Science*, pages 77–116, 1998. 136
- [Lod10] Florienne Loder. Taking a closer look inside blood vessels. http://www.mdhs.unimelb.edu.au/taking_a_closer_look_inside_blood_vessels, 2010. xv, 15
- [LT05] A. Lassó and E. Trucco. Vessel enhancement in digital x-ray angiographic sequences by temporal statistical learning. *Computerized Medical Imaging and Graphics*, 29(5):343–355, 2005. 38
- [LtOH10] Patrick J. Lynch, Fred the Oyster, and Mikael Haggstrom. coronary.pdf - this file is licensed under the creative commons attribution-share alike 3.0 unported license. http://en.wikipedia.org/wiki/File:Coronary_arteries.svg, 2010. xv, 3
- [Mac09] Carlo Maccia. Niveaux de référence et évaluation dosimétrique des patients. Les techniques interventionnelles en médecine et radioprotection - Société française de radioprotection, "http://www.sfrp.asso.fr/IMG/pdf/10-C_Maccia.pdf", oct 2009. 8
- [Mar87] D.M. Martinez. Model-based motion estimation and its application to restoration and interpolation of motion pictures. *NASA STI/Recon Technical Report N*, 88:13558, 1987. 32
- [MB06] Schoonenberg G Movassaghi B, Garcia JA. In vivo three-dimensional reconstruction of coronary stents based on motion compensated x-ray angiograms: first in human results. *Am J Cardiol*, 2006. 130
- [MCTTY10] Morice Marie-Claude, Hovasse Thomas, Lefevre Thierry, and Louvard Yves. Percutaneous coronary intervention for complex bifurcation lesions. *Interventional Cardiology*, 5:58–62, sept 2010. 76
- [Med11] St. Jude Medical. Pressurewire certus ffr measurement system copyright 2011 st. jude medical, inc. <http://www.sjmprofessional.com/Products/US/Intravascular-Sensors/PressureWire-Certus.aspx>, 2011. xvii, 27

- [Mei10] E. Meijering. Neuron tracing in perspective. Cytometry Part A, 77(7):693–704, 2010. [134](#), [135](#)
- [MF06] B. Mory and R. Florent. Multi-feature time filtering for enhancing structures in noisy images, June 1 2006. WO Patent WO/2006/056,924. [73](#)
- [MFKY11] Nick Mark, Paul Frake, Mia Kazanjian, and Lawrence Yu. Drug eluting stents. http://biomed.brown.edu/Courses/BI108/BI108_2004_Groups/Group05/Drug%20Eluting%20Stents/products_and_clinical_trials.htm, 2011. [8](#)
- [MJS⁺04] E. Meijering, M. Jacob, J.C.F. Sarria, P. Steiner, H. Hirling, and M. Unser. Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images. Cytometry Part A, 58(2):167–176, 2004. [134](#), [135](#), [136](#), [137](#)
- [MJSU03] E. Meijering, M. Jacob, JCF Sarria, and M. Unser. A novel approach to neurite tracing in fluorescence microscopy images. Signal and Image Processing, 399:96–148, 2003. [134](#), [135](#)
- [MNV99] E.H.W. Meijering, W.J. Niessen, and MA Viergever. Retrospective motion correction in digital subtraction angiography: a review. Medical Imaging, IEEE Transactions on, 18(1):2–21, 1999. [32](#), [84](#)
- [Mor09] Janine M Morris. 510k premarket notification submission: K092004. 2009. [29](#), [69](#)
- [MSB⁺10] S Milouchi, P Schiano, F Barbou, MC Chenilleau, and J Monsegu. Le rehaussement de contour de stent : un nouvel outil dans l optimisation du résultat de l rtangioplastie coronaire. 16eme assises du Collège National des Cardiologues des Hopitaux Généraux, (Posters), 2010. [75](#)
- [MSG⁺06] B. Movassaghi, D. Schaefer, M. Grass, V. Rasche, O. Wink, J. Garcia, J. Chen, J. Messenger, and J. Carroll. 3d reconstruction of coronary stents in vivo based on motion compensated x-ray angiograms. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2006, pages 177–184, 2006. [130](#)
- [MvdEV96] J.B.A. Maintz, P.A. van den Elsen, and M.A. Viergever. Evaluation of ridge seeking operators for multimodality medical image matching. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 18(4):353–365, 1996. [135](#), [136](#)

- [MVP⁺07] Jacob M. Mishell, Kalpesh T. Vakharia, Thomas A. Ports, Yerem Yeghiazarians, and Andrew D. Michaels. Determination of adequate coronary stent expansion using stentboost, a novel fluoroscopic image processing technique. Catheterization and Cardiovascular Interventions, 69(1), 2007. 31, 75
- [NIH10a] NIH. cabg.jpg - this image is in the public domain because it contains materials that originally came from the national institutes of health. http://www.nhlbi.nih.gov/health/dci/Diseases/cabg/cabg_what_is.html, 2010. xv, 7
- [NIH10b] NIH. Ptcanih.gif. this image is in the public domain because it contains materials that originally came from the national institutes of health. http://commons.wikimedia.org/wiki/File:PTCA_NIH.gif?uselang=fr, 2010. xv, 7
- [NIH10c] NIH. Ptcastentnih.gif. this image is in the public domain because it contains materials that originally came from the national institutes of health. http://fr.wikipedia.org/wiki/Fichier:PTCA_stent_NIH.gif, 2010. xv, 7
- [NIH10d] NIH. stentrestenosis02.jpg - what are the risks of coronary angioplasty? http://www.nhlbi.nih.gov/health/dci/Diseases/Angioplasty/Angioplasty_Risks.html, 2010. xv, 10
- [OG10] Boguslaw Obara and David Gavaghan Vicente Grau. Tracing curvilinear structures in live microscopy images. Symposium Microscopy Image Analysis for Biomedical Applications BMVA, London, UK, 2010. 134, 135, 136, 137
- [Org11] World Health Organization. Cardiovascular diseases (cvds), fact sheet n°317. <http://www.who.int/mediacentre/factsheets/fs317/en/index.html>, jan 2011. 3
- [OSB⁺08] Alice Ohanessian, Sanjay Sastry, Nabil Bendaoud, Thierry Lefevre, Thomas Hovasse, Yves Louvard, and Marie-Claude Morice. Abstract 4909: Usefulness of stent boost enhancement for stent expansion during coronary angioplasty in daily practice. Circulation, 118(18 Meeting Abstracts):S 958–, 2008. 31, 74, 75

- [OST93] M.K. Ozkan, M.I. Sezan, and A.M. Tekalp. Adaptive motion-compensated filtering of noisy image sequences. Circuits and Systems for Video Technology, IEEE Transactions on, 3(4):277–290, 1993. 32
- [Pau59] Hough Paul. Machine analysis of bubble chamber pictures. In Proc. of the Int. Conf. on High Energy Accelerators and Instrumentation (CERN), 1959. 135, 172
- [Pay01] M.M. Payne. Charles theodore dotter: The father of intervention. Texas Heart Institute Journal, 28(1):28, 2001. xv, 9
- [Per08] B. Perrenot. Reconstruction tomographique 3D de stent coronaire avec prise en compte du mouvement en angiographie rotationnelle cardiaque par rayons X. PhD thesis, INSA de Lyon, 2008. xxii, xxiii, 16, 103, 104, 106, 118, 119, 121, 128
- [ptc09] ptca.org. Medtronic’s endeavor drug-eluting coronary stent approved in japan. http://www.ptca.org/news/2009/0330_MEDTRONIC.html, 2009. xv, 10
- [PWBB97] D. Palti-Wasserman, A.M. Brukstein, and R.P. Beyar. Identifying and tracking a guide wire in the coronary arteries during angioplasty from x-ray images. BiomeA comparison of line enhancementtical Engineering, IEEE Transactions on, 44(2):152–164, 1997. 134, 135, 136, 137
- [RC08] Y. Rouchdy and L.D. Cohen. Image segmentation by geodesic voting. application to the extraction of tree structures from confocal microscope images. In Pattern Recognition, 2008. ICPR. 19th International Conference on, pages 1–5. IEEE, 2008. 134, 135, 159, 169
- [RC09] Y. Rouchdy and L.D. Cohen. The shading zone problem in geodesic voting and its solutions for the segmentation of tree structures. application to the segmentation of microglia extensions. In Proc. MMBIA 2009: IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis. IEEE, 2009. 169
- [RC11a] Y. Rouchdy and L.D. Cohen. A geodesic voting method for the segmentation of tubular tree and centerlines. In Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on, pages 979–983. IEEE, 2011. 169
- [RC11b] Y. Rouchdy and L.D. Cohen. A geodesic voting shape prior to constrain the level set evolution for the segmentation of tubular trees. In Proc.

- SSVM'11: 3rd International Conference on Scale Space and Variational Methods in Computer Vision, Ein Gedi, Israël (2011), 2011. [169](#)
- [RCT10] S. Redwood, N. Curzen, and M.R. Thomas. Oxford Textbook of Interventional Cardiology. Oxford Textbooks in Cardiology Series. Oxford University Press, 2010. [8](#)
- [RFF⁺07] Wayne Rosamond, Katherine Flegal, Gary Friday, Karen Furie, Alan Go, and et al. Heart Disease and Stroke Statistics–2007 Update: A Report From the American Heart Association Statistics Committee and Stroke Statistics Subcommittee. Circulation, 115(5):e69–171, 2007. [30](#)
- [RGLJA11] Veronique L. Roger, Alan S. Go, Donald M. Lloyd-Jones, and et al. Adams. Heart disease and stroke statistics–2011 update: A report from the american heart association. Circulation, 123(4):e18–209, 2011. [3](#)
- [RLH11a] Christopher Rohkohl, Günter Lauritsch, and Joachim Hornegger. Non-Periodic 3-D Motion Estimation and Reconstruction of Coronary Stents. In Marc Kachelrieß and Magdalena Rafecas, editors, Proceedings of 11th Fully 3D Meeting and 3rd HPIR Workshop, pages 462–465, 2011. [103](#)
- [RLH11b] Christopher Rohkohl, Gunter Lauritsch, and Joachim Hornegger. Non-Periodic 3-D Motion Estimation and Reconstruction of Coronary Stents. In Marc Kachelries and Magdalena Rafecas, editors, Proceedings of 11th Fully 3D Meeting and 3rd HPIR Workshop, pages 462–465, 2011. [131](#)
- [RLM⁺05] James C. Ross, David Langan, Ravi Manjeshwar, John Kaufhold, Joseph Manak, and David Wilson. Registration and integration for fluoroscopy device enhancement. In MICCAI, pages 851–858, 2005. [8](#), [15](#), [16](#), [31](#), [35](#), [36](#), [37](#), [38](#), [78](#)
- [RS89] A.R Rao and B.G. Schunck. Computing oriented texture fields. In Proceedings CVPR'8, 1989. [165](#)
- [SABM04] I. Stuke, T. Aach, E. Barth, and C. Mota. Estimation of multiple motions using block-matching and markov random fields. In SPIE Conf. Visual Communications and Image Processing VCIP2004, pages 274–278, 2004. [84](#), [86](#)

- [Sam86] R. Samy. An adaptive image sequence filtering scheme based on motion detection. Architectures and algorithms for digital image processing, pages 135–144, 1986. [32](#)
- [SB93] S. Sarkar and K.L. Boyer. Perceptual organization in computer vision: A review and a proposal for a classificatory structure. Systems, Man and Cybernetics, IEEE Transactions on, 23(2):382–399, 1993. [136](#), [137](#)
- [SBE⁺95] Gerhard Spekowius, H. Boerner, W. Eckenbach, Peter Quadflieg, and Gert J. Laurensen. Simulation of the imaging performance of x-ray image intensifier/tv camera chains. In Richard L. Van Metter; Jacob Beutel, editor, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, volume 2432, pages 12–23, May 1995. [36](#)
- [SF09] Gert Schoonenberg and Raoul Florent. Advanced visibility enhancement for stents and other devices: Image processing aspects. Cardiol Clin, 27:477490, 2009. [37](#), [38](#), [118](#), [130](#), [133](#), [189](#)
- [SFL⁺08] ”Gert Schoonenberg, Raoul Florent, Pierre Lelong, Onno Wink, Daniel Ruijters, and Bart ter Haar Romeny”. The effect of automated marker detection on in vivo volumetric stent reconstruction. MICCAI 2008, Part II, LNCS 5242, Springer-Verlag Berlin Heidelberg 2008, pages 87–94, 2008. [103](#), [104](#), [130](#)
- [SFL⁺09a] G. Schoonenberg, R. Florent, P. Lelong, O. Wink, D. Ruijters, J. Carroll, and B. Haar Romeny. Projection-based motion compensation and reconstruction of coronary segments and cardiac implantable devices using rotational x-ray angiography. Medical image analysis, 13(5):785–792, 2009. [16](#), [104](#), [130](#), [131](#)
- [SFL⁺09b] Gert Schoonenberg, Raoul Florent, Pierre Lelong, Onno Wink, Daniel Ruijters, John Carroll, and Bart ter Haar Romeny. Projection-based motion compensation and reconstruction of coronary segments and cardiac implantable devices using rotational x-ray angiography. Medical Image Analysis, 13(5):785 – 792, 2009. Includes Special Section on the 12th International Conference on Medical Imaging and Computer Assisted Intervention. [73](#)
- [SGH⁺04] H. Schmitt, M. Grass, S. Heiland, S. Hähnel, and K. Sartor. Image stacking with entropy values in conventional angiography: Initial experience1. Radiology, 230(1):294, 2004. [38](#)

- [SI04] B. Sarel and M. Irani. Separating transparent layers through layer information exchange. Computer Vision-ECCV 2004, pages 328–341, 2004. [84](#), [86](#), [87](#)
- [SI05] B. Sarel and M. Irani. Separating transparent layers of repetitive dynamic behaviors. In Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on, volume 1, pages 26–32. IEEE, 2005. [84](#), [86](#)
- [SJP⁺09] Ioannis Stathopoulos, Marcelo Jimenez, Georgia Panagopoulos, EUN Jung Kwak, Monica Losquadro, Howard Cohen, Sriram Iyer, Carlos Ruiz, Gary Roubin, and Kirk Garratt. The decline in pci complication rate: 2003-2006, versus 1999-2002. Hellenic J Cardiol, 50:379–387, 2009. [8](#)
- [SKD⁺99] P.W. Serruys, I.P. Kay, C. Disco, N.V. Deshpande, P.J. de Feyter, et al. Periprocedural quantitative coronary angiography after palmaz-schatz stent implantation predicts the restenosis rate at six months 1: Results of a meta-analysis of the belgian netherlands stent study (benestent) i, benestent ii pilot, benestent ii and music trials. Journal of the American College of Cardiology, 34(4):1067–1074, 1999. [8](#)
- [SKUF07] G. Slabaugh, K. Kong, G. Unal, and T. Fang. Variational guidewire tracking using phase congruency. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2007, pages 612–619, 2007. [134](#), [135](#), [137](#)
- [SL03] I.W. Selesnick and K.Y. Li. Video denoising using 2d and 3d dual-tree complex wavelet transforms. Wavelets: Applications in Signal and Image Processing X, 5207:607–618, 2003. [32](#)
- [SM91] M. Shizawa and K. Mase. Principle of superposition: A common computational framework for analysis of multiple motion. In Visual Motion, 1991., Proceedings of the IEEE Workshop on, pages 164–172. IEEE, 1991. [84](#), [86](#)
- [SMG07] D. Schafer, B. Movassaghi, and M. et al. Grass. Three dimensional reconstruction of coronary stents in vivo based on motion compensated x-ray angiography. Medical imaging 2007: visualization and image-guided procedures. Proceedings of SPIE, 6509, 2007. [130](#)
- [SNA⁺97] Yoshinobu Sato, Shin Nakajima, Hideki Atsumi, Thomas Koller, Guido Gerig, Shigeyuki Yoshida, and Ron Kikinis. 3d multi-scale line filter

- for segmentation and visualization of curvilinear structures in medical images. In CVRMed-MRCAS'97, volume 1205 of Lecture Notes in Computer Science, pages 213–222. 1997. [136](#)
- [SOF91] M.I. Sezan, M.K. Ozkan, and S.V. Fogel. Temporally adaptive filtering of noisy image sequences using a robust motion estimation algorithm. In Acoustics, Speech, and Signal Processing, 1991. ICASSP-91., 1991 International Conference on, pages 2429–2432. IEEE, 1991. [32](#)
- [Soi99] Pierre Soille. Morphological Image Analysis. Springer-Verlag, 1999. [48](#), [88](#)
- [SSD⁺05] G. Schoonenberg, M. Schrijver, Q. Duan, R. Kemkers, and A. Laine. Adaptive spatial-temporal filtering applied to x-ray fluoroscopy angiography. In Proc. SPIE, volume 5744, pages 870–878, 2005. [16](#)
- [ST98] P. Soille and H. Talbot. Image structure orientation using mathematical morphology. In Pattern Recognition, 1998. Proceedings. Fourteenth International Conference on, volume 2, pages 1467–1469. IEEE, 1998. [146](#)
- [Sti73] S.M. Stigler. The asymptotic distribution of the trimmed mean. The Annals of Statistics, pages 472–477, 1973. [40](#)
- [STR⁺05] V. Setlur, S. Takagi, R. Raskar, M. Gleicher, and B. Gooch. Automatic image retargeting. In Proceedings of the 4th international conference on Mobile and ubiquitous multimedia, pages 59–68. ACM, 2005. [91](#)
- [SvO07] J. Slotboom and D. van Ormondt. The effect of order-statistics filtering on the output probability density functions. Proc. 18th Annual Workshop on Circuits, Systems, and Signal Processing, pages 254–258, 2007. [39](#)
- [SVvdH⁺08] Tone Svilaas, Pieter J. Vlaar, Iwan C. van der Horst, Gilles F.H. Diercks, Bart J.G.L. de Smet, Ad F.M. van den Heuvel, Rutger L. Anthonio, Gillian A. Jessurun, Eng-Shiong Tan, Albert J.H. Suurmeijer, and Felix Zijlstra. Thrombus aspiration during primary percutaneous coronary intervention. New England Journal of Medicine, 358(6):557–567, 2008. [27](#)
- [Swe88] J.A. Swets. Measuring the accuracy of diagnostic systems. Science, 240(4857):1285, 1988. [139](#)

- [Swe96] J.A. Swets. Signal detection theory and ROC analysis in psychology and diagnostics: Collected papers. Lawrence Erlbaum Associates, Inc, 1996. 139
- [TA07] Hugues Talbot and Ben Appleton. Efficient complete and incomplete path openings and closings. Image Vision Comput., 25:416–425, April 2007. 135, 136, 146, 150, 151, 153
- [Tan10] Olena Tankyevych. Filtering of thin objects: applications to vascular image analysis. PhD thesis, Université Paris Est, 2010. 142, 143, 145, 146, 147, 172, 179
- [TDBP⁺09] Pim A.L. Tonino, Bernard De Bruyne, Nico H.J. Pijls, Uwe Siebert, Fumiaki Ikeno, Marcel van ‘t Veer, Volker Klauss, Ganesh Manoharan, Thomas Engstrøm, Keith G. Oldroyd, Peter N. Ver Lee, Philip A. McCarthy, and William F. Fearon. Fractional flow reserve versus angiography for guiding percutaneous coronary intervention. New England Journal of Medicine, 360(3):213–224, 2009. 27
- [tim05] NY times. Johnson to increase output of cypher stent. <http://www.nytimes.com/2005/08/18/business/18stent.html>, 2005. xvii, xx, 37, 83
- [TOM00] JG Toro, FJ Owens, and R. Medina. Multiple motion estimation and segmentation in transparency. In Acoustics, Speech, and Signal Processing, 2000. ICASSP’00. Proceedings. 2000 IEEE International Conference on, volume 6, pages 2087–2090. IEEE, 2000. 84, 86
- [TZPB01] D. Tubic, A. Zaccarin, J. Pouliot, and L. Beaulieu. Automated seed detection and three-dimensional reconstruction. i. seed localization from fluoroscopic images or radiographs. Medical physics, 28:2265, 2001. 88
- [VB10] Régis Vaillant and Vincent Bismuth. Procédé de traitement d’images radiologiques - french application number 1000076021, 2010. 18, 21, 82
- [VG05] Anthony J Viera and Joanne M Garrett. Understanding interobserver agreement: the kappa statistic. Family Medicine, 37(5):360–363, 2005. 62
- [vH92] Marcel van Herk. A fast algorithm for local minimum and maximum filters on rectangular and octagonal kernels. Pattern Recogn. Lett., 13(7):517–521, 1992. 48

- [Vin98] Luc Vincent. Minimal path algorithms for the robust detection of linear features in gray images. In Proceedings of the fourth international symposium on Mathematical morphology and its applications to image and signal processing, ISMM, pages 331–338, Norwell, MA, USA, 1998. Kluwer Academic Publishers. xxiv, xxvi, 135, 136, 149, 150, 151, 153, 155, 157, 158, 167, 169, 170
- [VLB06] Régis Vaillant, Jean Liénard, and Vincent Bismuth. System and method to enhance visualization of an object in a vascular vessel - us patent 7734328, 2006. 18, 20, 24, 29, 48
- [WA96] Y. Weiss and E.H. Adelson. A unified mixture framework for motion segmentation: Incorporating spatial coherence and estimating the number of models. In Computer Vision and Pattern Recognition, 1996. Proceedings CVPR'96, 1996 IEEE Computer Society Conference on, pages 321–326. IEEE, 1996. 84, 87
- [WCZ⁺09] P. Wang, T. Chen, Y. Zhu, W. Zhang, S.K. Zhou, and D. Comaniciu. Robust guidewire tracking in fluoroscopy. In Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, pages 691–698. IEEE, 2009. 135, 136
- [Wik11a] Wikipedia. Coronary circulation. http://en.wikipedia.org/wiki/Coronary_circulation, 2011. 3
- [Wik11b] Wikipedia. Intravascular ultrasound. http://en.wikipedia.org/wiki/Intravascular_ultrasound, 2011. 8
- [Wik11c] Wikipedia. Myocardial infarction. http://en.wikipedia.org/wiki/Myocardial_infarction, 2011. 4
- [Wik11d] Wikipedia. Robust statistics. http://en.wikipedia.org/wiki/Robust_statistics, 2011. 39
- [WSK⁺08] L. Wei, D. Shen, D. Kumar, R. Turlapati, and J.S. Suri. Effect of hierarchical deformable motion compensation on image enhancement for dsa acquired via c-arm. In Proceeding of the 20th Annual IS&T/SPIE Symposium on Electronic Imaging, pages 27–31, 2008. 32
- [WZL⁺07] J. Wang, X. Zhou, J. Lu, J. Lichtman, S.F. Chang, and STC Wong. Dynamic local tracing for 3d axon curvilinear structure detection from microscopic image stack. In Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium on, pages 81–84. IEEE, 2007. 134, 135, 136, 137

- [ZABT04] R. Zwiggelaar, S.M. Astley, C.R.M. Boggis, and C.J. Taylor. Linear structures in mammographic images: detection and classification. Medical Imaging, IEEE Transactions on, 23(9):1077–1086, 2004. [135](#), [142](#), [143](#), [145](#), [146](#), [147](#), [179](#)
- [ZAT⁺09] S. Zhang, C.K. Abbey, A. Teymoorian, X. Da, J.S. Whiting, and M.P. Eckstein. Motion and display effects on perception of multiple coronary stents. In Proceedings of SPIE, volume 7263, page 726306, 2009. [32](#)
- [ZHB⁺08] A.O. Zaid, I. Hadded, W. Belhaj, A. Bouallegue, S. Abdessalem, and R. Mechmeche. Improved localization of coronary stents based on image enhancement. International Journal of Biomedical Science, 4:212–216, 2008. [31](#), [74](#), [78](#)
- [ZLP⁺09] W. Zhang, H. Ling, S. Prummer, K. Zhou, M. Ostermeier, and D. Comaniciu. Coronary tree extraction using motion layer separation. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009, pages 116–123, 2009. [84](#), [86](#), [88](#)
- [ZPP04] V. Zlokolica, A. Pizurica, and W. Philips. Recursive temporal denoising and motion estimation of video. In Image Processing, 2004. ICIP'04. 2004 International Conference on, volume 3, pages 1465–1468. IEEE, 2004. [32](#)