



Curvilinear Structures Segmentation and Tracking in Interventional Imaging

Nicolas Honnorat

► To cite this version:

Nicolas Honnorat. Curvilinear Structures Segmentation and Tracking in Interventional Imaging. Other. Ecole Centrale Paris, 2013. English. NNT : 2013ECAP0007 . tel-00801865

HAL Id: tel-00801865

<https://theses.hal.science/tel-00801865>

Submitted on 18 Mar 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.



ÉCOLE CENTRALE DES ARTS
ET MANUFACTURES
<< ÉCOLE CENTRALE PARIS >>

Thèse
présentée par

Nicolas Honnorat

pour l'obtention du
GRADE DE DOCTEUR

Spécialité : Mathématiques Appliquées
Laboratoire d'accueil : MAS
SUJET :

**Segmentation et suivi de structures
curvilinéaires en imagerie interventionnelle**

soutenue le : 17 janvier 2013

devant un jury composé de:

Nassir NAVAB
Pierre-Dominique CROCHET
James S. DUNCAN
Alejandro FRANGI
Laurent NAJMAN
Wiro NIESSEN
Nikos Paragios
Régis Vaillant

Président
Examineurs

2013ECAP0007

ECOLE CENTRALE PARIS

P H D T H E S I S

to obtain the title of

Doctor of Ecole Centrale Paris

Specialty : Applied Mathematics

Defended by

Nicolas HONNORAT

Curvilinear Structures Segmentation and Tracking in Interventional Imaging

prepared at Ecole Centrale de Paris, CVC laboratory

defended on January 17, 2013

Jury :

<i>Reviewers :</i>	Alejandro FRANGI	- University of Sheffield, Sheffield
	Wiro NIESSEN	- Erasmus MC, Rotterdam
<i>Examiners :</i>	Pierre-Dominique CROCHET	- CHU de Nantes
	James S. DUNCAN	- Yale University
	Laurent NAJMAN	- Université Paris-Est
	Nassir NAVAB	- Technische Universität München
<i>Advisors :</i>	Nikos PARAGIOS	- Ecole Centrale Paris
	Régis VAILLANT	- General Electric Healthcare

Abstract

This thesis addresses the segmentation and the tracking of thin curvilinear structures. The proposed methodology is applied to the delineation and the tracking of the guide-wires that are used during cardiac angioplasty. During these interventions, cardiologists assess the displacement of the different devices with a real-time fluoroscopic imaging system. The obtained images are very noisy and, as a result, guide-wires are particularly challenging to segment and track.

The contributions of this thesis can be grouped into three parts. The first part is devoted to the detection of the guide-wires, the second part addresses their segmentation and the last part focuses on their spatio-temporal tracking.

Partial detection of guide-wires is addressed either through the selection of appropriate filter operators or using modern machine learning methods. First, a learning framework using an asymmetric Boosting algorithm for training a guide-wire detector is presented. A second method enhancing the output of a steerable filter by using an efficient tensor voting variant is then described.

In the second part, a bottom-up method is proposed, that consists in grouping points selected by the wire detector, in extracting primitives from these aggregates and in linking these primitives together. Two local grouping procedures are investigated: one based on unsupervised graph-based clustering followed by a line-segment extraction and one based on a graphical model formulation followed by a graph-based centerline extraction. Subsequently, two variants of linking methods are investigated: one is based on integer programming and one on a local search heuristic.

In the last part, registration methods are exploited for improving the segmentation via an image fusion method and then for tracking the wires. This latter is performed by a graph-based iconic tracking method coupled with a graph-based geometric tracking that encodes to certain extent a predictive model. This method uses a coupled graphical model that seeks both optimal position (segmentation) and spatio-temporal correspondences (tracking). The optimal solution of this graphical model simultaneously determines the guide-wire displacements and matches the landmarks that are extracted along it, what provides a robust estimation of the wire deformations with respect to large motion and noise.

keywords: MRF, Boosting, Steerable Filters, Tensor Voting, LP-stability based clustering, label costs, local search, iconic, geometric, detection, segmentation, tracking, registration

Résumé

Cette thèse traite de la segmentation et du suivi de structures curvilinéaires. La méthodologie proposée est appliquée à la segmentation et au suivi des guide-fils durant les interventions d'angioplastie. Pendant ces opérations, les cardiologues s'assurent que le positionnement des différents outils est correct au moyen d'un système d'imagerie fluoroscopique temps-réel. Les images obtenues sont très bruitées et les guides sont, en conséquence, particulièrement difficiles à segmenter.

Les contributions de cette thèse peuvent être regroupées en trois parties. La première est consacrée à la détection des guides, la seconde à leur segmentation et la dernière à leur suivi.

La détection partielle des guide-fils est réalisée soit par la sélection d'un opérateur de filtrage approprié soit en utilisant des méthodes modernes d'apprentissage artificiel. Dans un premier temps, un système réalisant un Boosting asymétrique pour entraîner un détecteur de guides est présenté. Par la suite, une méthode renforçant la réponse d'un filtre orientable au moyen d'une variante efficace de vote tensoriel est décrite.

Dans la seconde partie, une approche ascendante est proposée, qui consiste à regrouper des points sélectionnés par le détecteur de fil, à extraire des primitives des agrégats obtenus et à les lier. Deux procédures locales de regroupement des points sont étudiées : une reposant sur un clustering de graphe non supervisé suivi d'une extraction de segments de droites ; et l'autre reposant sur un modèle graphique puis une extraction d'axe central. Par la suite, deux méthodes de liaison des primitives sont étudiées : la première repose sur une approche de programmation linéaire, et la seconde sur une heuristique de recherche locale.

Dans la dernière partie, des méthodes de recalage sont utilisées pour améliorer la segmentation et pour suivre les fils. Le suivi proposé couple un suivi iconique avec un suivi géométrique contenant un modèle prédictif. Cette méthode utilise un modèle graphique déterminant à la fois une position du guide-fil (segmentation) et des correspondances (tracking). La solution optimale de ce modèle graphique décrit simultanément les déplacements du guide-fil et les appariements entre points d'intérêt qui en sont extraits, fournissant ainsi une estimation robuste des déformations du fil par rapport aux grands déplacements et au bruit.

mots-clef: MRF, Boosting, filtre orientable, vote tensoriel, clustering basé sur des stabilités LP, coût par label, recherche locale, iconique, géométrique, détection, segmentation, suivi, recalage

Acknowledgments

In this section, I would like to express all my gratitude to the people who have supported my research, that was partly conducted in Buc, in the GE Healthcare's facility and in Centrale Paris, in the team that is now known as the Center for Visual Computing.

First, I would like to thank my advisors, Nikos Paragios and Régis Vaillant, for their constant help and support to complete this Ph.D. thesis. Their supervision was very valuable for me. I thank them for having proposed this research topic, and organized this work, that was funded by the French National Technological Research Agency (ANRT) with the grant 1062/2008, by GE Healthcare and then by the French (INRIA). I thank in particular François Kotian that was my manager in GE Healthcare.

Then, I would like to thank my committee members: the Professors Wiro Niessen and Alejandro Frangi for reading and revising the draft of this thesis and for their insightful comments and the Professors J. S. Duncan, N. Navab, L. Najman and P.D. Crochet for taking part of the committee and for their interesting questions and valuable feedback.

I would like to thank Aristeidis Sotiras, Mickaël Savinaud, Stavros Alchatzidis and Chaohui Wang for their advice during the writing of the manuscript, and all the colleagues and the coworkers, that welcomed me in two very stimulating teams, both in Buc and in Centrale Paris: Jean Liénar, Cyril Riddell, Yves Trouset, Sebastien Gorges, Vincent Bismuth, Maxime Taron and Helene Langet in Buc ; and Chaohui Wang, Bo Xiang, Katerina Gkirtzou, Aristeidis Sotiras, Pierre-Yves Baudin, Mickaël Savinaud, Daniel Pescia, Stavros Alchatzidis, Stavros Tsogkas, Samuel Kadoury, Pascale Béliveau, Régis Behmo, Ahmed Besbes, Haithem Boussaïd, Iasonas Kokkinos, Martin de La Gorce, Vivien Fécamp, Fabrice Michel, Sarah Parisot, Loïc Simon, Olivier Teboul, Panagiotis Koutsourakis, Salma Essafi and Radhouène Neji in Centrale Paris, and in particular Carine Morotti-Delorme for her decisive help for organizing the defense.

I would like to thank James S. Duncan, that allowed me to visit his laboratory at Yale University during the summer 2010, where I had the opportunity to discuss with many people that I hope to meet again.

Last but not least, I would like to express all my gratitude to my family for its support during these years.

Contents

1	Introduction	10
1.1	Clinical Context	11
1.2	Purpose and Contributions	14
1.3	Related Fields of Research	15
1.4	Brief State-of-the-art	17
1.5	Thesis Outline	17
1.6	Publications	18
2	Curvilinear Structures Detection	19
2.1	Literature Review	20
2.1.1	Curvilinear Structures Detection	20
2.1.2	Curvilinear Structures Enhancement	31
2.1.3	Machine Learning Methods for Detection	35
2.1.4	State-of-the-art Guide-wire Detectors	37
2.2	Proposed Methods	38
2.2.1	Classifier Trained by Asymmetric Gentle AdaBoost	38
2.2.2	Steerable Filters Regularized by Tensor Voting	40
2.3	Experimental Validation	43
2.3.1	Validation Metrics	43
2.3.2	Testing Data	44
2.3.3	Guide-wire Detection Results with Boosted Classifier	44
2.3.4	Guide-wire Detection Results with Steerable Filters	45
2.4	Summary	46
2.5	Discussions	48
3	Curvilinear Structures Segmentation	50
3.1	State-of-the-art Methods	51
3.1.1	Tracing methods	53
3.1.2	Active Contours	59
3.1.3	Region Growing and Shortest Path Methods	60
3.1.4	Grouping Methods	63
3.1.5	Guide-wire Segmentation Methods	65
3.1.6	Conclusion	65
3.2	Contributions	66
3.2.1	Most Related Work	66

3.2.2	Line-segments Extraction	67
3.2.3	Elongated Structure Parts Extraction	71
3.2.4	Exploited Linking Criteria	73
3.2.5	Guide-wire Parts Linking	74
3.3	Experimental Validation	76
3.3.1	Experiments on Primitive Extraction	77
3.3.2	Experiments on Primitive Linking	79
3.4	Summary	85
3.5	Discussion	87
3.5.1	Limitations	87
3.6	Future Work	87
4	Registration for Wires Segmentation and Tracking	89
4.1	Introduction	90
4.2	State-of-the-art Methods	91
4.2.1	Geometric Tracking Methods	91
4.2.2	Iconic Tracking Methods	93
4.3	Linking Improvement via Graph Matching	95
4.4	Geometric-Iconic Graph-based Tracking	99
4.4.1	Iconic Tracking	101
4.4.2	Landmark Matching	102
4.4.3	Coupling Part	104
4.4.4	Dynamic Consistency	105
4.5	Experimental Validation	106
4.5.1	Segmentation through Graph/Image Fusion	106
4.5.2	Tracking	108
4.5.3	Segmentation and Tracking	110
4.5.4	Limitations of the Tip-distance Metric	112
4.6	Conclusions and Future Work	112
5	Review of Contributions and Future Work	118
5.1	Review of contributions	119
5.2	Review of Future Work	120
A	Demonstrations	122
A.1	Multiple Guide-wires Segmentation is NP-Hard	123

List of Figures

1.1	Fluoroscopic Images Particularities	12
1.2	Angioplasty intervention	13
1.3	Contributions of the thesis and related fields	16
2.1	Edge and ridge model	20
2.2	Tensor Voting schema	34
2.3	Assymetric boosting weights updates	40
2.4	ROC curve of two detector variants	46
2.5	Tensor-voting regularized steerable filters	47
2.6	Regularized steerable filter responses	47
3.1	HMM models	57
3.2	Elastica computation	63
3.3	Illustration of the adopted bottom-up pipeline	67
3.4	Ordering operations	76
3.5	Line-segments extraction	78
3.6	Microtubule delineation results	80
3.7	Long guide-wire parts extraction.	81
3.8	Comparison between the two primitive extraction methods	82
3.9	Line-segments extraction for the simple linking experiments	83
3.10	Quantitative detector performances	84
3.11	Quantitative simple linking results	84
3.12	Simple linking results	84
3.13	Linking criteria	85
3.14	Results provided by the ordering method	86
4.1	Line-segment matching	97
4.2	Tracking procedure	101
4.3	Quantitative performances of the complete method	107
4.4	Complete method for multiple GW delineation	108
4.5	Qualitative effects of the exploitation of the contrast-enhanced image	109
4.6	Sampling pattern and first tracking results	111
4.7	Missed Detection Rate over time	111
4.8	Segmentation results on the tracking database	113
4.9	Comparison of mean tracking results and tip distances	114

4.10 Performances obtained with or without semi-automatic initialization	115
A.1 Illustration of the proof	124

List of Tables

3.1	Line-segments extraction results	78
3.2	Quantitative results for the ordered linking on $D1$	85
4.1	Quantitative result of the complete ordering-based segmentation . .	107

Chapter 1

Introduction

1.1 Clinical Context: Cardiac Angioplasty Interventions

Myocardial infarction, also known as heart attack, is the leading cause of death in the developed countries. For instance, according to the World Health Organization [1], around 12.6% of deaths during 2002 were due to this disease (29.3% were due to cardiovascular diseases in general). According to the recent survey [2], 7 900 000 Americans suffer (or have suffered) from heart attack in 2012 while the coronary heart diseases represent 16 300 000 patients.

Myocardial infarctions are due to the accumulation of lipids (cholesterol and fatty acids) inside the blood vessels feeding the cardiac muscle, the coronary arteries, that leads to the formation of atherosclerotic plaques locally reducing the radius of the blood vessels (stenosis). The sudden rupture of these plaques delivers solid pieces that restrain or block the blood flow in a downstream portion of the coronary artery. The blood flow restriction provokes the partial destruction of the cardiac muscle and leads, without intervention, to the death of the patient.

Myocardial infarctions can be treated in two distinct manners. The first refers to a heart bypass that consists in creating new branchings in the coronary artery in order to guide blood flow through a path that does not include the affected part of the artery. This kind of interventions requires full surgery with direct access to the heart of the patient. For this reason, they are complicated and dangerous for the him.

A different family of methods, the *angioplasty methods*, has been introduced in 60' following the pioneering work of Dotter [3], that is considered as *the father of Interventional Radiology* and won the Nobel prize in 1978. Such procedures are micro-invasive methods and they consist of pushing a tube, called *catheter*, into the blood vessels of the patient until it reaches the artery to be cured. This initial step, called *catheterization*, can benefit from potentially different entry points (arm, in the radial artery or in the femoral artery) and was developed in the 30' by Werner Forssmann, that earn the Nobel prize in 1956. Once the catheter is correctly placed, a guide-wire is then pushed into this pipe, and used as a support to position a metallic armature called stent enveloping a deflated balloon at the location of the stenosis. The inflation of the balloon simultaneously reshapes the blood vessels and deploys the stent that is let there for reinforcing the blood vessel wall. This intervention is illustrated by the figure [Fig. (1.2)].

During the intervention, the cardiologist assesses the position of the surgical devices through a real time X-ray imaging system, also known as *fluoroscopic* imaging device. Fluoroscopy exists since the end of the nineteenth century [4]. Since X-rays are harmful, these devices have been designed to employ small X-rays doses. As a result, the shot noise leading to a signal-to-noise ratio proportional to the square root of the number of photons received by each X-ray detector unit, the acquired signal exhibits a low signal-to-noise ratio.

The majority of existing fluoroscopic systems use one planar 2D X-rays detector (also called *flat-panel detector*), and acquire images of size 1000×1000 pixels




Figure 1.1: (a) raw fluoroscopic images. (b) compressed images, before and during contrast agent injection. Catheters are indicated in yellow, guide wires tips in green and deployed stents blue. Red arrows indicate confusing structures that are sometimes visible in clinical images: the rib borders.




Figure 1.2: Angioplasty intervention with stent deployment (NIH image)

or less, at a frame rate taken between 10 and 30 images per second. These images are compressed by an algorithm that preserves the contrasts that are useful for the cardiologist and confer the image a nice visual appearance. Since the compression is non-linear and the algorithm an important marketing element that is kept secret by the fluoroscopic device manufacturers, the compressed images are no more described by the Beer-Lambert-Bouguer law and the width of the observed structures can no more be related to the contrast observed in the image formed with the logarithm of the observed X-rays intensity. Moreover, in spite of the high spatial resolution (around 0.2 millimeter per pixel) and the high temporal resolution (less than 0.1 seconds between two successive images) the guide-wires are extremely difficult to detect since their width is approximately 2-3 pixels large. The guide-wires undergo also a large non-rigid motion (often higher than 25 pixels between successive frames) resulting from the combination of cardiac and respiratory motion. Furthermore, the nature of observations (2D data corresponding to the projection of 3D wire) leads to self-crossing structures as well as to self/auto auto-occlusions. The clinical context further complexifies the process due to the presence of non-relevant but still similar thin and curvilinear structures. Last but not least, the guide-wire is made of two parts that look and move very distinctly: the *body* made of a flexible and moderately transparent material and a *tip* that contains a radio-opaque and more flexible material. All these characteristics, that are illustrated in the figure [Fig. (1.1)] make the segmentation of the guide-wires very challenging for an automatic method.

Such a task would however be of crucial interest, since it would allow to develop methods helping the cardiologist during the interventions. In particular, it would allow to develop interesting interactions such as automatic measurement of the width of the artery of interest, automatic checking of the deployed stent, etc. The segmentation of stents was addressed in the recent work [5]. We refer the reader to the thesis [6] for more insight about this task and its relation with the segmentation of the guide-wires, that attracts our interest in this work.

1.2 Purpose and Contributions

The purpose of this thesis is to segment the guide-wires that are used during cardiac angioplasty interventions. Since this task is difficult and time consuming, and since efficient tracking algorithms were proposed for wire-like structures, we decided to segment the guide-wire in the first frame of the fluoroscopic sequences only, and then to track it in order to describe its deformations from frame to frame. Since both segmentation and tracking require the extraction of the support of the wire in the fluoroscopic images, we have studied this task separately. We have thus finally addressed three tasks: the detection of the guide-wire, the segmentation of the guide-wire and its tracking.

In this manuscript, two different methods are considered for the guide-wire detection. The first relies on machine learning techniques, while the second exploits a dedicated family of steerable filters for ridge detection and a tensor voting method.

A complete bottom-up pipeline for guide-wire segmentation is then proposed, that selects pixels according to the detection output, groups them and extracts primitives from the clusters obtained (either a line segment representing the group or a polygonal line approximating its centerline) and finally links together these primitives in order to fully reconstruct the wire.

Finally, a tracking approach combining an iconic parametrized curve tracking method with a geometric/landmark based tracking method into an hybrid graph-based framework is presented.

1.3 Related Fields of Research

A curvilinear structures detector is necessary for extracting the support of the guide-wires. Such detectors date from the early years of computer vision with particular focus on the extraction of object borders. Numerous methods have been proposed and are currently used in a very broad range of domains, such as remote sensing and medical image analysis. These methods have been employed to address various image analysis tasks such as biological structures delineation (blood vessels in X-ray or retinal fundus images [7, 8, 9], microtubules and actin filaments in Total Internal Reflection Fluorescence images [10, 11], neurites [12, 13], etc) and roads segmentation in aerial and satellite images [14, 15, 16].

The performances of the structure detection is crucial, since the outcome is exploited directly as a data term by the subsequent segmentation frameworks. Our work adopts this principle as well, where the output of the detection is either exploited for providing a set of points likely to belong to the wire (considered as input to our segmentation method) or directly exploited as a data term (by our tracking method). We also considered curvilinear structures enhancement methods during this work, because the low signal-to-noise ratio of the fluoroscopic images and the width of the wires make the detection of the guide-wires extremely challenging.

Almost all the segmentation methods extracting roads in aerial or blood vessels/neurites in medical images are applicable to our problem. Moreover, the bottom-up approach that we propose for segmenting the guide-wires involves the same kind of ideas as surface reconstruction methods and multi-part objects segmentation frameworks. These methods require the definition of linking criteria and are prone to perform a global optimization.

Last but not least, many contour tracking methodologies could be applied to our tracking problem, and strong connections exist between the registration literature and the wire tracking method that we developed. Our tracking approach can indeed be considered as a frame-to-frame registration of the wire, and rely on a combination of 1D variants of existing 3D registration methods.

The figure [Fig. (1.3)] illustrates the basis of the schemes that are described in this work.




Figure 1.3: The three boxes describe the main contributions of this work and the alternatives that are proposed: for instance, a guide-wire detection module based on asymmetric Boosting is investigated, as well as an alternative relying on steerable filters and tensor voting in the Chapter 2. The fields of research related to the different methods are indicated outside the boxes.

1.4 Succinct State-of-the-art on Guide-wires Segmentation and Tracking

Early work on the segmentation and the tracking of catheters, that are larger and more opaque than guide-wires, dates from to the beginning of the 90'. Amini et al. have thus expressed the tracking of catheters into a probabilistic formulation, a Markov Random Field, solved by dynamic programming in [20].

Numerous approaches have been proposed since then for the segmentation of the guide-wires. One can classify them into three groups: tracing methods [21, 22], shortest path methods [23] and bottom-up/grouping methods such as [24]. The guide-wire is detected in various manner: by using Hessian filter [25] as in [26], by computing the Vesselness [27] like in [21], by the mean of phase congruency [28] like in [29], or machine learning based detector, such as in [24] that exploits the Probabilistic Boosting Tree method [30]. Many approaches, originated from the curve tracking literature, have also been proposed for the tracking of the wires. One can group them into two families: snake-splines [26, 29] and probabilistic based tracking [21, 31, 32, 33].

All these methods are examined in detail in the corresponding parts of this manuscript. Since many alternatives exist for each task (detection, segmentation and tracking of the wires) the related fields of research are described extensively before we present our contributions. A particular attention is given to the approaches developed in the context of the medical or biological imaging.

1.5 Thesis Outline

The remainder of the thesis is organized as follows.

The chapter 2 reviews the main approaches applicable for the extraction of the visual support of the guide-wire and describes our contribution to that domain. The support extracted is used for the segmentation as well as for the tracking, and its quality is critical for both of them.

The chapter 3, after a review of the most important curvilinear structures segmentation methods, presents a modular pipeline performing guide-wires segmentation. Different variants of the main modules of the framework are investigated in detail.

The chapter 4 describes the methods that were used so far for addressing the tracking of curvilinear structures and the method that we have developed for the purpose of guide-wire tracking.

A review of the contributions and a discussion conclude the manuscript.

1.6 Publications

This work rely on the following publications:

- N. Honnorat, R. Vaillant and N. Paragios. Robust Guide-wire Segmentation Through Boosting, Clustering and Linear Programming. In IEEE ISBI, pages 924927, 2010
- N. Honnorat, R. Vaillant and N. Paragios. Guide-wire Extraction through Perceptual Organization of Local Segments in Fluoroscopic Images. In MICCAI, 2010
- N. Honnorat, R. Vaillant, James S. Duncan and N. Paragios. Curvilinear Structures Extraction in Cluttered Bioimaging Data with Discrete Optimization Methods. In ISBI, 2011
- N. Honnorat, R. Vaillant and N. Paragios. Graph-based Geometric-Iconic Guide-Wire Tracking. In MICCAI, 2011
- N. Honnorat, R. Vaillant and N. Paragios. Graph-based Guide-wire Segmentation through Fusion of Contrast-enhanced and fluoroscopic Images. In IEEE ISBI, 2012

Chapter 2

Curvilinear Structures Detection




Figure 2.1: (i) edge model (ii) Gaussian ridge model (iii) an edge detector detects these two profiles approximately in the same manner.

Thin, curvilinear structures appear in many contexts. Their detection is of prominent interest for many tasks and has been the subject of a considerable amount of interest. In this chapter, we first review the existing approaches for curvilinear structures detection. We present then our contributions and their experimental validation. Discussions conclude the chapter.

2.1 Literature Review

In this section, we first review the literature dealing with curvilinear structures detection. We present then curvilinear structures enhancement methods -that are used to significantly improve the detections- and the machine learning approaches that can be used to train high performance detectors. We finally focus on the work that was specifically carried out on guide-wires detection.

2.1.1 Curvilinear Structures Detection

There exist many different kinds of boundary discontinuities: edges, ridges, borders between textured regions... In order to provide a comprehensive review, we will restrict our discussion on methods targeting ridges detectors because the guide-wires appear as dark ridges and edges detectors because they can indirectly serve to detect ridges, as shown in [Fig. (2.1)]. In the sequel, we review the main curvilinear detectors that could be applied for detecting the guide-wires that are used during cardiac angioplasty. We have classified these methods into the following categories:

- **differential/integral approaches** can be divided into **intensity based methods**, that regroup all the filters and all the operators based directly on the image intensity or on its derivatives and **flux-based methods** relying on the computation of the flow of a vector field that has been extracted from the image in order to detect the structures of interest.

- **mathematical morphology methods** exploit two basic operators (dilation and erosion) to enhance boundary discontinuities.
- **quadrature filters** methods convolve images with pairs of quadrature filters. They consider the two responses obtained as the two parts of a complex number and compute a local phase and a local norm. Gabor filters constitute a famous example. In this section, we focus first on anisotropic wavelets, that were developed to overcome the limitations of wavelets towards proper handling of 2D discontinuities [34, 35]. In most of the approaches, the norm is the only part that is exploited. A family of methods exploiting only the phase was however proposed [28]. By considering weighted phase spread, they provide measures that are less sensitive to contrast variations. We will present these methods in detail.
- **steerable filters** are efficient operators that can provide boundary strength and orientation. Contrary to classical filters, the response of these filters can be obtained for any orientation - by combining the responses obtained when convolving the image with a set of basis filters.
- **problem-specific operators**, are developed according to their application. They exploit the properties of the problem at hand in order to solve it efficiently and accurately.

Differential/Integral Approaches

The first curvilinear detectors that were developed are based on the convolution of the image with derivative and smoothing filters. Hereafter, we will only present the continuous smoothed operators that are commonly employed. These variants constitute an improvement with respect to the related non-smoothed operators, because the smoothing (by a Gaussian function) changes the scale at which the image is considered and allows to discard high frequency noise. In order to reduce the computational time, each operator has moreover been discretized, leading to a broad variety of derivating operators. The well known Sobel operator is for instance a discretized non-smoothed gradient operator.

In this section, we will denote with $I(x, y)$ an image, with $I_x(x, y), I_y(x, y)$ its first derivatives and $I_{xx}(x, y), I_{xy}(x, y), I_{yy}(x, y)$ its second derivatives.

Image Intensity based filters

The Difference of Gaussian (DoG) is one of the simplest filters that can be used to select structures according to their spatial width. It consist in convolving the image with two Gaussians of different scales, and subtract the result obtained with the largest scale from the other. Because the Fourier transform of a Gaussian function is a Gaussian function, and because the Fourier transform is linear, convolving an image with a DoG is equivalent to multiplying its Fourier transform with a DoG, that is a bandpass operation. DoG is also used to approximate the

Laplacian of Gaussian (LoG), for instance when extracting SIFT interest points [36, 37].

Structure Tensor

The structure tensor [38] is based on the image intensity gradients. Let $w(.,.)$ be a positive windowing function (for instance, a 2D Gaussian), the structure tensor is given by:

$$\mathcal{T}(x, y) = \sum_{p, q} w(p - x, q - y) \begin{pmatrix} I_x(p, q)^2 & I_x(p, q)I_y(p, q) \\ I_x(p, q)I_y(p, q) & I_y(p, q)^2 \end{pmatrix}$$

Diagonalizing this tensor allows to determine the main orientation of the gradient in the neighborhood of the current point, and to differentiate the edges (where the two eigenvalues are very different) from the blobs (where the two eigenvalues are equal) and the flat regions (where they are equal and small).

Because the structure tensor is given by a weighted average over a neighborhood of the considered location, it is more reliable than the original gradient. The same kind of computation can be done with other vector fields, as in [39], where the structure tensor is used as a regularization scheme, and as we do in section [Sec. (4.4.2)].

Hessian based Methods

The second order derivative can be used to determine the location of the ridges, because perfect dark ridges correspond to a minimum of the image intensity in the direction orthogonal to them and a constant image intensity in the direction parallel to them. Let $I(x, y)$ be an image, the matrix $H(x, y)$ of the second order derivatives of I is called Hessian:

$$H(x, y) = \begin{pmatrix} I_{xx}(x, y) & I_{xy}(x, y) \\ I_{xy}(x, y) & I_{yy}(x, y) \end{pmatrix}$$

Diagonalizing this matrix provides the local normal to the ridge (that is given by the eigenvector related with the highest eigenvalue) and the sharpness of the ridge (that is related to the value of this eigenvalue).

In practice, in order to extract more reliable ridges and ridges at different scales, the convolution of the image with different Gaussian functions is considered. Deriving the convolution of an image with a filter and convolving an image with the derivative of the filter are equivalent. Thus, if g_σ is the Gaussian function of zero mean and standard deviation σ , and $g_{\sigma,xx}, g_{\sigma,xy}, g_{\sigma,yy}$ denote its second order derivatives, then the Hessians at scale σ , H_σ are given by:

$$H_\sigma = \begin{pmatrix} g_{\sigma,xx} * I & g_{\sigma,xy} * I \\ g_{\sigma,xy} * I & g_{\sigma,yy} * I \end{pmatrix}$$

The Laplacian of Gaussian (LoG) filter (also known as Marr–Hildreth filter [40]) is quite similar:

$$LoG = g_{\sigma,xx} * I + g_{\sigma,yy} * I$$

It is used for detecting edges, by exploiting the property that the LoG is equal to zero at the location of the edge.

Several measures detecting blood vessels - known as *Vesselness* measures for this reason - and based on the Hessian eigenvalues have been proposed. We describe three of them, denoting with $|\lambda_1| < |\lambda_2|$ the two Hessian eigenvalues:

- the **Lorenz vesselness** [41] is given by:

$$\mathcal{I} = \sigma^\gamma \frac{|\lambda_2|}{|\lambda_1|}$$

where $\gamma > 0$ is a parameter that is used to compare the responses at different scales and is determined according to the intensity profile that is sought (Gaussian profile, bar like profile, etc).

- the **Sato's vesselness** [42] is defined in 3D for bright tubular structures as:

$$\mathcal{I} = \begin{cases} |\lambda_1| \left(\frac{\lambda_2}{\lambda_1} \right)^{\gamma_{12}} \left(1 + \frac{\lambda_3}{|\lambda_2|} \right)^{\gamma_{23}} & \lambda_1 < \lambda_2 < \lambda_3 \leq 0 \\ |\lambda_1| \left(\frac{\lambda_2}{\lambda_1} \right)^{\gamma_{12}} \left(1 - \alpha \frac{\lambda_3}{|\lambda_2|} \right)^{\gamma_{23}} & \lambda_1 < \lambda_2 < 0 < \lambda_3 < \frac{|\lambda_2|}{\alpha} \\ 0 & otherwise \end{cases}$$

where $\gamma_{12}, \gamma_{23} \geq 0$ control the sharpness of the isotropy required for the cross section of the tubular structure, and $0 < \alpha \leq 1$ induces asymmetric results with respect to the sign of the highest eigenvalue λ_3 .

- the **Frangi's vesselness** [27] is probably the most famous. It is given (for dark ridges in 2D images) by:

$$\mathcal{I} = \begin{cases} 0 & if \lambda_2 < 0 \\ e^{-\frac{R_B^2}{2\beta^2}} \left(1 - e^{-\frac{S^2}{2\gamma^2}} \right) & otherwise \end{cases}$$

where $R_B = \lambda_1/\lambda_2$ is a blobness measure that is small for narrow elongated structures (and tends to 1 for isotropic blobs). $S^2 = \lambda_1^2 + \lambda_2^2$ is used to discard regions where the image intensity variations are small.

Some filters were directly designed on Gaussian derivatives. In [43], a filter with increased orientation selectivity has been obtained by summing more than two shifted Gaussian derivative filters (summing two Gaussian derivatives only approximates the Hessian filter).

Flux-based Methods

Flux-based methods rely on the integral over a surface of interest of a vector field extracted from the image. For instance, if gradients are extracted from the image, an algorithm seeking the deformable surface traversed by the maximal flux will provide a segmentation that relies on strong edges. Such a method was used for detecting curvilinear structures centerlines in [44].

When trying to segment a vector field, taking the orientation of the vector into account as well as their norm is obviously of prominent interest. This is the reason why the authors of [45] proposed a flux-maximizing level set method. During the evolution described by their method, the active surface tries to maximize the flux through its border. As demonstrated in their article, such an idea is particularly appealing for segmenting blood vessels, because the gradient norm is high at the wall of the blood vessels. [46] describes an acceleration of this method based on Fourier Transform. It is however less clear if this method could handle very thin structures like the guide-wires (that are sometimes only 2-3 pixels broad).

In order to use smooth vector fields (and make them more reliable for snakes methods, see chapter 4), the gradient vector flow [47] was introduced. The prominent interest of this framework resides in the fact that it produces a smooth vector field by diffusing an edge map.

Let f be an edge map, i.e. a function that has highest values near the edge of the structures of interest. The GVF framework determines the vector field $w(x, y) = (u(x, y), v(x, y))$ that minimizes (where μ is a positive constant):

$$E = \int \mu(u_x^2 + u_y^2 + v_x^2 + v_y^2) + |\nabla f|^2 |w - \nabla f|^2 dx dy$$

This leads to a smooth vector field wherever the edge map is constant and a vector field approximating the edge map gradient where it is strong. The solution is obtained by solving the associated Euler equation with an iterative scheme.

Because this scheme diffuses the edge map information through the image, GVF provides a reliable information for snake frameworks. Nowadays it is intensively used for tubular structures segmentation such as blood vessel segmentation in [48] and neurites in [49].

GVF has been used in most of the applications that were exploiting gradient vectors. It has been used as basis for a criterion similar to the Frangi's vesselness [27]. In [50, 48], the local Jacobian of the GVF is used instead of the classical Hessian matrix to describe a mono scale, smoothed, vesselness measure.

Mathematical Morphology based Methods

In this section, we briefly review the main mathematical morphology operators [51] that can be used to extract curvilinear structures.

The two basic mathematical morphology operators are the erosion and the dilation by a structuring element. These operators modify the image intensity of each pixel by taking into account the information of a neighborhood described

by a structure element. Let I and S denote a gray-level image and a structuring element, that is a set of coordinates, then the dilation $\oplus$ and the erosion $\ominus$ are defined by:

$$\begin{aligned} [I \oplus S](x) &= \max_{(z-x) \in S} I(z) \\ [I \ominus S](x) &= \min_{(z-x) \in S} I(z) \end{aligned}$$

These two operators do not commute, and are combined to create new operators: an *opening* is obtained by applying an erosion followed by a dilation. Its name comes from the fact that this operation discards thin bright structures, and can therefore open the holes that are enclosed in a bright structure. Similarly, a dilation followed by an erosion is called a *closing* because it can fill the holes of a structure. Last but not least, the *white top-hat* operator is the identity minus an opening, and the *black top-hat* a closing minus the identity. The first extracts structures that are smaller than the structuring element and brighter than their surrounding, while the latter extracts small structures that are darker than their neighborhood. This last operator can be used to extract guide-wires as explained in [52].

Recent advances in the field of mathematical morphology include the development of two distinct categories of methods for boundary detection. First, spatially variant morphology operators were developed. These operators process pixels according to local information such as Vesselness measures [53, 54]. Second, a mathematical definition of the paths has been proposed, and morphological operators on them were presented and efficiently implemented [55].

Quadrature Filters

A pair of quadrature filters f, g is a pair of filters such that g is the Hilbert transform of f . Their responses are considered as the real and imaginary parts of a complex number. This combined signal is known as the analytic signal.

Anisotropic Wavelets

The wavelet representations are a powerful tool in signal processing theory. By allowing multiscale processing, wavelets provide robust and sparse representations of natural signals, that are useful for a broad variety of applications (signal compression, signal denoising ...).

Several wavelets were developed to address the detection of lines and curves.

- **Beamlets** [56] can be considered as a multiscale set of appropriately located line-segments. These wavelets are particularly interesting for extracting contours from noisy backgrounds.

- **Ridgelets** were proposed and studied in [34]. They were slightly modified in [35] in order to constitute an orthonormal set of functions. They rely on integrals over continuous sets of ridges orientations. An orthogonal ridgelet is thus indexed by its scale, its location, its angular scale, its angular location and a gender token.

One can also cite **curvelets** [57, 58], that are based on ridgelets, **contourlets** [59, 60] and **directionlets** [61] that were developed as an alternative relying on separable filtering operations. Many of these wavelets have been used in medical image analysis. In [62], for instance, the contourlets are used to enhance the blood vessels visible in the retinal fundus images.

Phase based Methods

Instead of exploiting the modulus of the analytic signal, some methods rely exclusively on the phase for detecting edges. The principle underlying this approach is that the phase is locally congruent at the location of the edges - and that the phase spreads are significant. One can thus detect an edge by summing the phase spreads at different scales. This is done in [28], where an approach for diminishing the influence of phase noise and the influence of small spreads are also presented. Their **phase congruency** $\mathcal{PC}(x)$ at location x in an image is based on oriented wavelet filters:

$$\mathcal{PC}(x) = \frac{\sum_o \sum_n W_o(x) (A_{on}(x) \Delta \Phi_{on}(x) - T_o)_+}{\sum_o \sum_n A_{on}(x) + \epsilon}$$

where ϵ is a small constant, o denotes the orientation of the filter exploited (for instance Gabor filters in [28]), n accounts for the different scales, $A_{on}(x)$ is the amplitude obtained at scale n and orientation o , $\phi_{on}(x)$ the phase ($\bar{\phi}_o(x)$ the average of the $\phi_{on}(x)$ over n), T_o is calculated for canceling the effect of the noise, the $(\cdot)_+$ function avoids negative corrected phase spreads:

$$(x)_+ = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases}$$

In [28], the spread function is given by:

$$\Delta \Phi_{on}(x) = \cos(\phi_{on}(x) - \bar{\phi}_o(x)) - |\sin(\phi_{on}(x) - \bar{\phi}_o(x))|$$

Last but not least, the weighting functions $W_o(x)$ is used to discard locations where the spread is small.

The main advantage of the phase congruency (and the phase-based methods) is its ability to detect low contrasted curvilinear structures (because the extracted measures are less sensitive to contrast variations). This is the reason why it was used for detecting guide-wires in [29]. The main disadvantage comes from the fact that low contrasted confusing or noisy structures are detected as strongly as the sought ones. Although this issue has been partially addressed in [28], it remains a limitation that has to be taken into account.

Steerable Filters

Steerable functions have been introduced in [63]. In 2D, a steerable function f written in cylindrical coordinates (ρ, θ) is a function exhibiting a finite Fourier decomposition:

$$f(\rho, \theta) = \sum_{k \in K \subset \mathbb{Z}} a_k(\rho) e^{jk\theta}$$

Properties

Let f be a steerable function, then rotating f by an angle ϕ yields:

$$f(\rho, \theta + \phi) = \sum_{k \in K \subset \mathbb{Z}} a_k(\rho) e^{jk\theta} e^{jk\phi}$$

For this reason, convolving an image I with a steerable filter f_ϕ rotated by ϕ yields:

$$I * f_\phi = \sum_{k \in K \subset \mathbb{Z}} g_k(\phi) (I * f_k)$$

where $f_k(\rho, \theta) = a_k(\rho) e^{jk\theta}$ and $g_k(\phi) = e^{jk\phi}$.

In other words, the convolution of an image with any rotated version of a steerable filter is generated by a linear combination of the responses obtained when the image is convolved with a finite set of basis filters. The analytical expression of the $g_k(\phi)$ being known as soon as the Fourier decomposition of the filter has been described, the steerable filters are used in the following manner:

- the image is convolved with the basis filters f_k , in order to obtain the $I_k = (I * f_k)$.
- the orientation yielding the maximal filter response is given by:

$$\phi^* = \operatorname{argmax}_\phi \sum_{k \in K \subset \mathbb{Z}} g_k(\phi) I_k = \operatorname{argmax}_\phi g(\phi)$$

- ϕ^* and the corresponding response constitute the output of the filter.

When the optimization problem is easy to solve (when the polynomial in $\tan(\phi)$ obtained for $g(\phi)$ is of degree lower than four), most of the computational burden is due to the convolution between the image and the basis filters. In order to accelerate the computations, many authors [63, 64, 65] prefer using x-y separable filters:

$$f(x, y) = \sum_{l, m} h_{l, m}(\theta) x^l y^m$$

It is straightforward to express these filters in the generic form (based on Fourier decomposition) when observing that:

$$\begin{aligned} x &= \rho \cos(\theta) = \rho \frac{e^{j\theta} + e^{-j\theta}}{2} \\ y &= \rho \sin(\theta) = \rho \frac{e^{j\theta} - e^{-j\theta}}{2j} \end{aligned}$$

The Gaussian derivatives is an other important class of steerable filters. Since the derivatives in any direction are obtained by linearly combining the derivatives with respect to x and y of the same order and since these derivatives are x-y separable, all the Gaussian derivatives are efficiently computable steerable filters. The well-known Hessian filter is a famous example.

The analytical formula for rotating a linear combination of Gaussian derivatives was presented in [64]. For the sake of completeness (and because we present and use steerable filters based on Gaussian derivatives in the next sections) we recall this formula. Let f be a steerable filter combining the derivatives of a Gaussian function $g(x, y)$ in the following manner:

$$f(x, y) = \sum_{k=1}^M \sum_{i=0}^k \alpha_{k,i} \left(\frac{\partial^{k-i}}{\partial x^{k-i}} \frac{\partial^i}{\partial y^i} g(x, y) \right) \triangleq \sum_{k=1}^M \sum_{i=0}^k \alpha_{k,i} g_{k,i}(x, y)$$

then rotating f by an angle θ lead to the rotated filter:

$$f^\theta(x, y) = \sum_{k=1}^M \sum_{i=0}^k \beta_{k,i}(\theta) g_{k,i}(x, y)$$

where:

$$\begin{aligned} \beta_{k,i}(\theta) &= \sum_{j=0}^k \sum_{l,m \in S(k,j,i)} \binom{k-j}{l} \binom{j}{m} (-1)^m (\cos\theta)^{j+l-m} (\sin\theta)^{k-j-l+m} \\ S(k, j, i) &= \{l, m | 0 < l < k-j; 0 < m < j; k-(l+m) = i\} \end{aligned}$$

Steerable Filters for Ridges Detection

In [63], contour detection is addressed by means of the *orientation energy*. By definition, if $G_n(\theta)$ is the n^{th} derivative of the Gaussian filter along the first coordinate rotated by the angle θ , the orientation energy $E_n(\theta)$ is defined as:

$$E_n(\theta) = [G_n(\theta)]^2 + [H_n(\theta)]^2$$

$H_n(\theta)$ is the Hilbert transform of $G_n(\theta)$ (in other words: $G_n(\theta)$ and $H_n(\theta)$ form a quadrature pair of filters). The orientation energy is approximated by:

$$E_n(\theta) \approx C_1 + C_2 \cos(2\theta) + C_3 \sin(2\theta)$$

The local orientation, θ_d , is given by the phase of this approximation:

$$\theta_d = \frac{1}{2} \text{Arg}(C_2 + jC_3)$$

and the phase of the quadrature pair (the argument of the complex number: $G_n(\theta_d) + jH_n(\theta_d)$) is used to distinguish the different features.

For instance, in order to detect contour, Freeman and Adelson propose to use the second order oriented energy $E_2(\theta)$ [63]. They characterize the pixels belonging to a contour by the fact that their oriented energy is maximum in the directions orthogonal to their local orientation θ_d . Similarly to Canny [66] detector, a thresholding with hysteresis is then applied. The phase allows to distinguish edges from ridges, dark and white lines.

Steerable Filters for Ridges Detection designed with Canny-like Criteria

In [64], special families of edges, ridges and wedges detectors based on Gaussian derivatives are derived from criteria inspired from the criteria used by Canny to design his famous edge detector [66]. These criteria are expressed under an integral form and assembled into a quadratic criterion, whose minimization with a Lagrange multiplier method leads to an eigenvalue problem.

More precisely, when denoting with $\mathbf{g}$ the set of Gaussian derivatives with respect to the coordinates x and y , the filter $h(x, y)$ is designed by:

$$h = \mathbf{a}^T \mathbf{g}$$

where $\mathbf{a}$ is a vector of parameters to determine.

Let us assume that the feature to detect is described by the function $f_0(x, y)$ when it is aligned with the X-axis (horizontal orientation) and let denote with $\mathbf{g}_i$ the derivative of Gaussian number i . Let us introduce the inner product (between function being part of $\mathcal{L}^2$):

$$\langle f, g \rangle = \int_{\mathbb{R}^2} f(x, y)g(x, y)dx dy$$

The three criteria introduced in [64] are:

- **the signal-to-noise-ratio:** the signal S provided by the filter is:

$$S = \int_{\mathbb{R}^2} f_0(x, y)h(-x, -y)dx dy = \langle f_0, -h \rangle$$

If the signal is corrupted by additive white noise of unit variance, then the variance of the noise, N , is equal to:

$$N = \int_{\mathbb{R}^2} |h(x, y)|^2 dx dy = \langle h, h \rangle$$

Like Canny edge detector [66], a good filter is a filter that maximizes the ratio $\frac{S^2}{N}$.

- **the localization:** a good filter produces a response that is maximal at $y = 0$ and as sharply maximal as possible. Therefore, maximizing the following value ensures that the response decreases quickly around the peak at $y = 0$:

$$L = - \int_{\mathbb{R}^2} f_0(x, y) h_{yy}(-x, -y) dx dy = \langle f_0, h_{yy} \rangle$$

- **the smoothness:** a good detector does not produce oscillations with high amplitude (parallel to the X-axis). In order to reduce the bandwidth of the filter, the authors [64] remark that it is sufficient to penalize:

$$R_o = \int_{\mathbb{R}^2} |h_{yy}(x, y)|^2 dx dy = \langle h_{yy}, h_{yy} \rangle$$

to reduce the oscillation orthogonal to the feature, and to penalize:

$$R_p = \int_{\mathbb{R}^2} |h_{xx}(x, y)|^2 dx dy = \langle h_{xx}, h_{xx} \rangle$$

to smooth the response along the feature, yielding to the penalty:

$$R = R_o + R_p.$$

Let μ denote a positive constant, Jacob and Unser propose to maximize the criterion [64]:

$$C = SL - \mu R$$

under the constraint: $N = 1$ (the filter has been normalized).

Replacing h by its definition in terms of linear combination of Gaussian derivatives leads to express the signal S , the localization L , the oscillation R and the criterion N in the following manner:

$$\begin{array}{ll} S = \mathbf{a}^T s & s_i = \langle f_0, -\mathbf{g}_i \rangle \\ L = \mathbf{a}^T l & l_i = \langle f_0, (\mathbf{g}_i)_{yy} \rangle \\ R = \mathbf{a}^T V \mathbf{a} & \text{with } V_{i,j} = \langle (\mathbf{g}_i)_{yy}, (\mathbf{g}_j)_{yy} \rangle + \langle (\mathbf{g}_i)_{xx}, (\mathbf{g}_j)_{xx} \rangle \\ N = \mathbf{a}^T P \mathbf{a} & P_{i,j} = \langle \mathbf{g}_i, \mathbf{g}_j \rangle \end{array}$$

As a result, C is a quadratic function of $\mathbf{a}$:

$$\begin{aligned} C &= \mathbf{a}^T [U - \mu V] \mathbf{a} \\ \text{where: } U &= s l^T \end{aligned}$$

Introducing a Lagrange multiplier λ in order to take the constraint: $N = \mathbf{a}^T P \mathbf{a} = 1$ into account and taking the derivative of the obtained expression leads to:

$$P^{-1} (U - \mu V) \mathbf{a} = -\lambda \mathbf{a}$$

The best $\mathbf{a}$ is thus the eigenvector $\mathbf{a}_\lambda$ normalized such that $N = 1$ and C is maximized. In [64], several filters are designed in this manner to detect ridges, edges and wedges. In particular, a second-order and a forth-order ridge detector (involving respectively second and forth order derivatives of Gaussian) are presented. Using these filters requires to solve a quadratic (resp. quartic) polynomial in $\tan(\theta)$ to find the optimal orientation, which can be done analytically.

Problem-specific Operators

In the context of blood vessels segmentation in retinal fundus images, for instance, these filters are called **matched filters** and were proposed in [67]. A matched filter simply contains the template of an ideal blood vessel. This template is rotated and scaled as many time as required by the user, and the image to filter is convolved with every resulting filter. Matched filters have been widely used for retinal blood vessels segmentation, and for microtubules segmentations [68]. [52] is an another example of filters devoted to the enhancement of long ridges. This filter is defined for eight orientations, and takes integers values in order to be efficiently computed (this filter can be related to the kernels used in [69, 70] to extract edges).

The main limitation of such methods rely on the fact that the rotated and scaled filters are not separable and not sparse, which results in computational inefficiency. This concern is however tackled by the use of simplified filters/kernels and the recent availability of massively parallel and cheap calculus architectures: the programmable graphic cards.

Remaining detectors comprise simple or ad hoc approaches, as for instance in [71], where the structure orientation is given by a Principal Component Analysis of the local gradients.

2.1.2 Curvilinear Structures Enhancement

Until now we have presented local detection methods. Unfortunately, these methods fail to cope with less visible or occluded parts of the structure of interest.

Several methods have been introduced for addressing this issue by propagating the detection score from the regions where they are strong to the rest of the image, in order to fill the gaps between these reliable regions. In this section, we first present the main families of anisotropic diffusion schemes [72] that were used for enhance curvilinear structures. We present then the tensor voting scheme [73].

Anisotropic Diffusion

The equation describing heat diffusion has been used since the 80' for diffusing the intensity level through noisy images in order to smooth them. Let $I_0(x)$ be an image. This model leads to the following PDE:

$$\begin{aligned} I(x; 0) &= I_0 \\ \frac{\partial I(x; t)}{\partial t} &= \Delta I(x; t) = \text{div}(\nabla I(x; t)) \end{aligned}$$

A very important variant of **anisotropic diffusion** is the **Mean Curvature Motion** (MCM) [74]:

$$\frac{\partial I(x; t)}{\partial t} = |\nabla I(x; t)| \text{div} \left(\frac{\nabla I(x; t)}{|\nabla I(x; t)|} \right)$$

By applying the classical property:

$$\operatorname{div}(f\vec{A}) = f\operatorname{div}(\vec{A}) + \nabla f \cdot \vec{A}$$

to the previous equation, the Laplacian part orthogonal to the gradient is the only that remains. Thus, MCM only produces diffusion along level lines.

An interesting extension of this framework consists in locally modifying the diffusion according to a tensorial field $D(x; t)$ [75, 76, 77]:

$$\frac{\partial I(x; t)}{\partial t} = \operatorname{div}(D(x; t)\nabla I(x; t))$$

where $D(x; t)$ is a tensor (i.e. a symmetric definite positive matrix) that is built from the image. For instance, a variant called **coherence enhancing diffusion** [78] relies on the structure tensor $\nabla I(x; t)\nabla I(x; t)^T$. The purpose of this approach is to reinforce the diffusion parallel to the edges and reduce the diffusion in the orthogonal direction. The Mean Curvature Motion can be modified in the same manner.

Shock filters [79, 80, 81] are an alternative to anisotropic diffusion. These filters are used to restore edges and therefore can be used to reinforce the detection of the two boundaries of a tubular structure.

Anisotropic diffusion is explained in detail in the following book [72]. We refer also the reader to the review [82], that presents vessel dedicated anisotropic diffusion scheme, and we point out the work [83] that specifically focuses on the enhancement of blood vessels (and exploits a variant of the Frangi Vesselness [27] to construct the diffusion tensors). The recent work [84] compares different diffusion scheme (for liver blood vessel enhancement).

Tensor Voting

Tensor voting [85, 73, 86] is a well-known method for structure enhancement in tensorial images.

A tensor is a symmetric positive semi-definite real matrix. It is thus a diagonalizable matrix with positive eigenvalues. The eigenvectors are orthogonal to each other. In 2D, one can therefore completely characterize a tensor by its eigenvalues and the orientation θ of the eigenvector related to the highest eigenvalue - called main axis. A tensor can also be decomposed into a sum of a **ball-tensor** without orientation containing only the information of the smallest eigenvalue and a **stick-tensor** - that we will also call **stick-vector** - that encodes the eccentricity of the tensor and the orientation of its main axis. Let λ_1 and λ_2 be the eigenvalues of a tensor $\mathcal{T}$, with $\lambda_1 > \lambda_2$, the **ballness** b of the tensor is equal to λ_2 , its **stickness** is the difference between the two eigenvalues: $s = \lambda_1 - \lambda_2$ and its **orientation** is the orientation of the main eigenvector. During 2D tensor voting, the tensors are described by their ballness, stickness and orientation. If e_1 and e_2 are the (unit)

eigenvectors associated with λ_1 and λ_2 , $\mathcal{T}$ is given by:

$$\mathcal{T} = \lambda_1(e_1 \otimes e_1) + \lambda_2(e_2 \otimes e_2) = (\lambda_1 - \lambda_2)(e_1 \otimes e_1) + \lambda_2 \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Tensor voting consists in making each location of the tensor field send votes in its neighborhood. The votes are tensors that depend on the location where they were emitted and on the relation between the emitter and the receiver. The output of the tensor voting is obtained by summing of all votes. In the case of stick-voting, the input tensor have a null ballness, and each location votes for another place at most once. A ball-voting, by contrast, is expressed as a sum of stick-votes, and the expression for ball-voting results from integration calculus.

The stick-votes (votes transmitted by stick-tensors) are proportional to the eigenvalue of the stick tensor, and oriented according to its orientation. One can thus obtain any stick-vote field from **the vote cast by the tensor** $\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$, that is called the **voting field**.

The voting field was defined in [85, 73] according to following criteria:

- the vote emitted by a stick-tensor is a stick-tensor
- the **co-circularity** criterion constrains the emitting stick-vector and the received stick-vector to be orthogonal to the same circle passing through the emitting and the receiving locations.
- the exponential decay along the circle: the norm of the received vote decreases exponentially with the distance along the circle described previously.
- the exponential decay with the curvature of the circle described previously.

Let denote with O the vote emitter (where the stick-vector is parallel to the x-axis and unitary) and A the location described with (r, ϕ) in cylindrical coordinates receiving the vote. The following voting field was introduced:

$$V(r, \phi) = e^{-\frac{s^2 + ck^2}{\sigma^2}} \begin{pmatrix} 1 - \cos(4\phi) & -\sin(4\phi) \\ -\sin(4\phi) & 1 + \cos(4\phi) \end{pmatrix}$$

Where s is the length of the arc of circle between O and A , k the curvature of this circle, c a parameter weighting the curvature penalization and σ is a scale. Figure [Fig. (2.2)] presents the decomposition of the tensors, the stick-votes and these values. Applying simple trigonometric formulas leads to:

$$s = \frac{r\phi}{\sin(\phi)}$$

$$k = \frac{2\sin(\phi)}{r}$$

The ball vote is given by an integration of infinitesimal stick votes over the possible directions.




Figure 2.2: (i) tensor decomposition (ii) stick vote (red), received vote (blue) and lengths and angles used to describe the voting field

Interpretation/Applications

One can interpret stick-votes as support with respect to the normal of the structure of interest.

Since the seminal work [85], tensor voting has been applied to various problems of increasing complexity and dimension, such as curve completion in 2D images, plane completion, stereo-vision, dimensionality estimation, manifold learning [87, 88] and function approximation [88] among others. A detailed description of the original tensor voting and its main application can be found in [87]. In the biomedical imaging field for example, tensor voting is used in [89] to improve the segmentation of neuron membranes in electron microscopy image stacks as a post-processing method.

Relations between tensor voting and the stochastic completion fields

A contour completion method inspired by the principles of the Gestalt's psychology [90, 91, 92] and relying on random walks was proposed in [93]. This method, called *stochastic completion field* is further described in [94]. It represents the likelihood that a particle sent from one fragment of the contour and received by an other one passes through a given location with a given orientation.

As explained in [95], there exist strong similarities between the stochastic completion fields and the tensor voting: both cast a directional vote, either in one step - tensor voting - or through a Markovian diffusion process - stochastic completion fields-. In [95], a relation between a tensor and an orientation bundle is derived. This relation is exploited towards a tensor voting variant approximating the stochastic completion fields, that is compared with the standard tensor voting [85].

The main limitation of tensor voting resides in the high computational burden resulting from the rotation of the voting field. Several approaches have been investigated to tackle this drawback. In [85] the input tensor field was constrained to be sparse. Such an approach is obviously not applicable when the input vector field is dense and when enhancing small structures - that are discarded by the

thresholding making the tensor field sparse. In [96] the exponential decay function depending on the curvature, the distance and the angle is replaced by a new decay function that can be determined using a closed form formula. In [97], the exponential decay is replaced by a steerable function that is efficiently rotated. Another approach, adopted in [98], consists in decomposing the different influences of the different votes (stick, plate and ball voting in 3D) and in simplifying the votes resulting on significant computational complexity reduction.

Geodesic Voting

Geodesic voting relies on a similar principle: a set of geodesic paths is determined, and then each path vote for the point that it visits. In this manner, the points that belong to interesting locations where many paths overlap are enhanced, contrary to the points belonging to short isolated structures.

In [99], shortest paths from a seed point given to the algorithm are used. Because these paths are computed from an action map obtained from a curvilinear structure detector, they tend to pass through curvilinear structures, that allows to enhance the portions of these structures that are less detected. In [100], weighted polygonal paths arriving at every point of the image are used for the same task.

2.1.3 Machine Learning Methods for Detection

In this section, we briefly present some important machine learning tools and their applications to contour/curvilinear structures detection.

Machine learning addresses the following tasks (among others):

- **reinforcement learning** determines the best decision to take according to a set of rewards provided by the environment in return to these actions.
- **supervised binary classification** consists in learning, from a labeled training data, a function correctly predicting the labels of so-far-unseen *testing* data. Given a training set (x_i, y_i) of samples $x_i \in \mathcal{X}$ labeled as *positive* ($y_i = +1$) or *negative* ($y_i = -1$), a function $f : \mathcal{X} \mapsto \{-1, +1\}$ is learned. This function is called *classifier* because it distinguishes the samples of the two (positive and negative) classes.
- **regression** can be seen as an extension of the classification to real labels. Regression aims at constructing/determining a function $f : \mathcal{X} \mapsto \mathbb{R}$ predicting the value y_i associated to the x_i .
- **ranking** is an extension of the classification that seeks a *scoring* function that can be mapped by a monotonic increasing bijection to the labels. In other words, this scoring function should rank the data correctly.
- **unsupervised classification** consists in recognizing a structure underlying the data, made of a set of classes, and building the corresponding classifier. Like other unsupervised methods, it requires no labels.

Numerous problematics have emerged, such as the ability to deal with multiple classes (multi class classification), the ability to deal with uneven training sets (where some classes are rare), the ability to deal with missing data labels (denoted as *semi-supervised* learning) or with partial annotations of group of related data (multiple instance learning).

We focus on the methods that are the most appropriate for our problem: the detection of the pixels that belong to the guide-wires. We will thus restrict ourselves to the supervised classification methods, considering the pixels part of the wires as belonging to a positive class, and the background of the fluoroscopic images as the negative class.

The Support Vector Machines (SVM) are among the most famous methods. They determine the separating hyperplane lying as far as possible from the classes samples, while separating them correctly. The smallest (Euclidean) distance between this hyperplane and the samples is called the **margin**. By dividing all the problem parameters by this margin, this maximization problem becomes a minimization of the hyperplane normal under the assumption that all the points have a margin greater than 1. Due to the fact that a hyperplane perfectly separating two classes is often impossible to determine, slack variables penalizing the difference between the actual margin of the points and the unit-margin are introduced and the dual optimization problem is solved. The solution of the problem is expressed under the form of a weighted sum of *support vectors*, that are actually the data that are the nearest to the hyperplane. A regression version of the SVM was developed in [101]. Support Vector Machines have been already intensively used to address the detection of curvilinear structures. In [102, 103] SVMs exploiting Gabor filters were introduced to detect blood vessels in retinal fundus images. In [65], a rotation invariant detector based on steerable functions (Gaussian derivatives) is built with SVM. This detector is used again in [104, 105, 106].

Kernel methods rely on the use of positive semidefinite kernels, called Mercer kernels (kernels being in that context a continuous symmetric function of two variables). According to Mercer's theorem, these kernels can be considered as inner products in a space of higher dimension where the data would have been projected by an appropriate function. The kernel trick [107] consists in replacing inner products (into machine learning methods) by kernel functions, such as to project the data in a space of higher dimension where they will be easier to handle. For instance, in [108], the data is projected in a high dimension space where the separation with an hyperplane is feasible and more precise compared to the one that could have been achieved in the original space. This method is particularly attractive because it is neither necessary to define the function used for the projection nor useful to compute the projection: choosing a kernel is sufficient, and many simple kernel functions can be used, such as the exponential and the polynomial kernels. The kernel trick is now commonly applied in the case of SVMs, in order to obtain non-linear separating surfaces.

The Boosting methods [109, 110] progressively build a strong classifier by linearly combining weak classifiers. Equal weights are attributed to the training

samples. Sequentially, a weak classifier minimizing an error depending on the weights is chosen, it is added to the existing ones (with the correct coefficient) and the weights are changed so as to focus the training on the data that are misclassified during the next steps. Many variants were proposed improving the pioneering work of [109]. Modification of the error function of Boosting can produce variants for regression and ranking [111, 112]. A statistical description of the Boosting is developed in [110]. We refer the readers to the review [113] for a more detailed analysis of the Boosting methods.

The probabilistic boosting tree [30] (PBT) is a variant of boosting algorithm that builds a cascading tree [114]. It has been applied to train contour detectors like for example in [115], where the binary version of the PBT is used (the positive class corresponding to points belonging to the contours, and the negative one to the remaining part of the images). The same idea was amended to provide also an estimation of the local orientation in [24].

There exist many other classification methods. We will only briefly describe two of them:

- **the Fisher’s linear discriminant** determines an hyperplane separating two classes. By contrast to the SVM, this hyperplane is the one that maximizes the separation criterion obtained by dividing the inter-classes variance by the intra-classes variance. The hyperplane is finally found by solving an eigenvalue problem related to this criterion.
- A **Decision tree** is a tree containing one criterion per interior node, that determines a choice between the child nodes. The leaves are associated with classes. Thus, a data is classified by following the path determined by the value of the criteria encountered from the root to a leaf, that provides the most likely label for this data. Decision trees are naturally able to perform multi-class classification. Random forests [116] are made of multiple trees whose criteria have been (partly) randomly chosen. Data is classified by using all the trees of the forest and keeping the label obtained by the dominant vote. We refer the reader to [117] for more details on splitting criteria used for decision trees.

In the recent work [118], linear filters detecting curvilinear structures are determined by a dictionary learning method. The output of these filters, together with an Optimally Oriented Flux [119, 46] and the Frangi’s Vesselness [27], are used as input for a Random Forest [116] or a logistic regression [120] classifiers.

2.1.4 State-of-the-art Guide-wire Detectors

Most of the families of curvilinear structure detectors have already been applied to the detection of guide-wires. In [121], a modified Laplacian and a modified Marr-Hildreth filters are used to detect the wire. In [26, 122], the Hessian is used. The vesselness [27] is employed in [21, 32] as support for tracking the guide-wire. The phase congruency [28] has been exploited in [29]. A boosted detector exploiting

the multiclass version of probabilistic boosting tree [30] has been used in [24] to detect edgels belonging to the guide-wires

2.2 Proposed Methods

In this section, we present two new guide-wires detection methods, respectively based on asymmetric gentle AdaBoost and on tensor voting and steerable filters.

2.2.1 Classifier Trained by Asymmetric Gentle AdaBoost

Let recall that the Boosting algorithm [109] constructs a binary classifier exhibiting high performance (therefore called *strong* classifier) by linearly combining *weak* classifiers, that perform just better than random classifiers. The method is iterative and determines at each step one weak classifier and its corresponding weight by minimizing a weighted sum of errors over the training set. Each time a new weak classifier is chosen and added to the strong one, the weights of the training data are modified so as to put more focus on the data that are still misclassified (and less on the data that are well classified).

If $H(\cdot)$ denotes the strong classifier constructed, and $(X, Y) = (x_i, y_i)_{i=1 \dots N}$ denotes the training data, where $y_i \in \{-1, 1\}$, the Boosting algorithms try to minimize the empirical estimation of the following error:

$$\begin{aligned} J(H(x)) &= \mathbb{E} [e^{-y(H(x))}|x] \\ &\approx \frac{1}{N} \sum_{i=1}^N e^{-y_i(H(x_i))} \end{aligned}$$

Let denote with $h(\cdot)$ a weak classifier. The variant called *Gentle Adaboost* in [110] adopts a Newton stepping procedure to minimize the error. This algorithm is described in figure [Alg. (1)].

Algorithm 1 Gentle Adaboost

given data set $(x_i, y_i), i = 1 \dots N$, M number of weak learners

initialize: $w_i = 1/N$

For $m = 1 \dots M$

 learn weak classifier h_m using w_i by fitting x to y

$H(x) \leftarrow H(x) + h_m(x)$

$w_i \leftarrow w_i \exp(-y_i h_m(x_i))$

output model: $G(x) = \text{sgn}(H(x))$

This algorithm is stable, because for each x , the weak classifier h finally chosen verifies:

$$h(x) = P_w(y = 1|x) - P_w(y = -1|x) \in [-1, 1]$$

It performs therefore often more robustly than other boosting variants like, for instance, the *Real AdaBoost* variant that, at each iteration, minimizes exactly the energy $J(H(x) + h(x))$ over $h(x)$ [110]. This is the reason why we adopted this variant.

For guide-wires detection, however, the data is extremely asymmetric. The positive class - the pixels of the wires - are proportional to the size of the image while the background is proportional to its number of pixel. This asymmetry induces an exclusive focus of the learning process on the background pixels. In order to still detect the positive class, addressing this challenge is necessary. Three different approaches can be envisaged to do so.

First of all, it is possible to reduce the dominant class by only retaining a subset of the training data of the same size as the small one. When this subset is randomly sampled with replacement, this approach is said to exploit *bootstrapping*. This idea is attractive for cascade learning [114, 123], because the misclassified data are gradually reduced. On the contrary, when a unique classifier is trained, one need to select the most confusing data from the dominant class. In [124], multiple trainings are thus performed by using all the data of the small class and a growing subset of the other class. This second set contains at the beginning a small random subset of the data, that is completed after each training by a random subset of the misclassified data.

One can also **bootstrap the small class**: this methods is based on augmenting the data set through sampling randomly from the rare class (with replacement) toward increasing the size of this cohort without violating its statistical properties. As described in the review [125], under-sampling strategies like the previous ones and over-sampling strategies like this one are numerous in the data mining literature.

Besides, several dedicated methods handling imbalanced training sets have been developed. We have adopted the **asymmetric Boosting** method [126] that puts more focus on the rare class by gradually increasing the weights of its samples and less emphasis on the other class by decreasing the weights of its samples. Other Boosting variants could have been applied as well, such as the method presented in [127] that modifies less the samples weights of the dominant class with respect to the classical Boosting update. **Adacost** [128] takes into account a cost depending on the samples when updating their weights (and constitutes therefore a generalization of the previous methods). Last but not least, **RareBoost** [129] employs different weights update rules for the different classes. Because Adacost does not consider separately the four kinds of samples (true or false, positive or negative), it may perform poorer than RareBoost that considers all of them (as explained in [125]). The weights update of AdaBoost and the aforementioned algorithms are presented in the figure [Fig. (2.3)]. We refer the interested readers to the review [125] for more details and other methods (for instance, methods addressing data fragmentation and inappropriate bias).

Method	weights update	
Adaboost	$w_i \leftarrow w_i \exp(-\alpha_m y_i h_m(x_i))$	
Asymmetric Adaboost	$w_i \leftarrow w_i \exp(-y_i(\alpha_m h_m(x_i) - \beta))$	β positive constant parameter
[127]	$w_i \leftarrow \begin{cases} w_i \exp(-\alpha_m y_i h_m(x_i)) & \text{if } y_i = 1 \\ w_i \exp(-\alpha_m y_i h_m(x_i)\nu) & \text{else} \end{cases}$	$\nu < 1$ positive constant parameter
AdaCost	$w_i \leftarrow w_i \exp(-\alpha_m y_i h_m(x_i)\beta(i))$	$\beta(i) = f(\text{sgn}(y_i h_m(x_i)), c_i)$ where c_i are the cost

Figure 2.3: classical AdaBoost weights update and weight updates designed to address uneven training sets. The rare class is by convention the positive one. The function f (empirically chosen) in the AdaCost update allows to put more or less focus on the different samples according to the costs c_i chosen by the user before starting the training.

We could not increase the size of the training set and we did not want to perform a selection among the negative samples. Therefore, we employed the Asymmetric AdaBoost update [126] in our Gentle AdaBoost method. We obtained a stable AdaBoost algorithm adapted to our dramatically unbalanced training problem. In order to obtain a classifier output that can be transformed into a likelihood, we translated the output by $-M\beta$, using finally the following classifier (still with the notations 2.1.3):

$$H(x) = \left(\sum_{m=1 \dots M} h_m(x) \right) - M\beta$$

The algorithm employed by our method is finally the one described in [Fig. (2)]. During the validation, the parameter β was determined by cross-validation.

2.2.2 Steerable Filters Regularized by Tensor Voting

The main limitation of the supervised Machine Learning methods is that they require a training, and that the trained detector is unable to adapt to completely different data. When new data is used, one has to train a new detector. In the context of clinical fluoroscopy, many different kinds of images sequences are handled such as raw images provided directly by the X-ray detectors, images processed by lossy compression, record images with high contrast and poor contrast interventional images. It is too complicated to train a new detector per application. The complexity of this approach is also too high (without parallelization) for real-time application.

Aiming to develop a generic detector exploitable for the tracking of guide-wires, we have decided to use a new detector. This detector exploits only one kind of steerable filter for ridge detection [64] and enhances its response by tensor voting

Algorithm 2 Asymmetric Gentle Adaboost

given data set $(x_i, y_i), i = 1 \dots N$, asymmetry parameter β , M number of weak learners

initialize: $w_i = 1/N$

For $m = 1 \dots M$

learn weak classifier h_m using w_i by fitting x to y

$H(x) \leftarrow H(x) + h_m(x)$

$w_i \leftarrow w_i \exp(-y_i(h_m(x_i) - \beta))$ and renormalize

output model: $\hat{G}(x) = \text{sgn}(H(x) - M\beta) = \text{sgn}\left(\left(\sum_{m=1 \dots M} h_m(x)\right) - M\beta\right)$

real output: $G(x) = \left(\sum_{m=1 \dots M} h_m(x)\right) - M\beta$

[85, 86]. Noting that the performances improvement introduced by the forth-order filter was minor compared to the increase in computational burden, we used the second-order filter exhibiting the highest signal-to-noise [64]. This filter was independently designed for neurite tracing in [130].

Let us denote with $v_\theta(x, y)$ the guide-wire orientation estimated by the steerable filter and $r(x, y)$ the response of the filter, both considered at location (x, y) . Our method exploits the following tensors:

$$\mathcal{T}(x, y) = r(x, y) [v_\theta(x, y) \otimes v_\theta(x, y)]$$

where $\otimes$ denotes the tensor product. These tensors are then processed by the fast tensor voting scheme [97]. The stickness and the orientation of the stick tensors of the tensors obtained constitute the guide-wire detection and the local guide-wire orientation estimated by our method.

Let us describe briefly the method [97]. As explained in section [Sec. (2.1.2)], tensor voting consists in making each location of a tensorial image vote in its vicinity and add these votes. A tensor of the form :

$$C\mathcal{T}_\alpha = Cu_\alpha \otimes u_\alpha$$

where u_α is the vector of unit norm in direction α , casts votes that are C times larger than the rotation by α of the votes sent by $\mathcal{T}_0$. These latter constitute the **voting field**.

As pointed out in [97], the complexity of the Tensor Voting originates from the rotations of the voting field. In [97] the exact voting field designed in [85, 86] is replaced by an approximated Voting Field exhibiting a limited Fourier decomposition when expressed in cylindrical coordinates. Such a voting field is steerable and can be efficiently rotated.

More precisely, let V denote a steerable voting field, V_{xx}, V_{yy}, V_{xy} its components, V^θ its rotation by θ , $U(x, y)$ the final tensorial image, $s(x, y)$ the stickness at location (x, y) , $\theta(x, y)$ the orientation of the stick-tensor and $\mathcal{I}$ the entire image.

Then:

$$U(x, y) = \int_{\mathcal{I}} s(h, k) V^{\theta(h, k)}(x - h, y - k) dh dk$$

Recalling the following property of the 2D rotation group:

$$\begin{aligned} \vec{V}^\theta &= S^{-1} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{-2j\theta} & 0 \\ 0 & 0 & e^{2j\theta} \end{pmatrix} S \vec{V} \\ \text{with } S &= \begin{pmatrix} 1 & 0 & 1 \\ 1 & -2j & -1 \\ 1 & 2j & -1 \end{pmatrix} \\ \text{and } \vec{V} &= \begin{pmatrix} V_{xx} \\ V_{xy} \\ V_{yy} \end{pmatrix} \end{aligned}$$

the authors perform the rotations on the following vectors:

$$\begin{pmatrix} V_0 \\ V_2 \\ V_{-2} \end{pmatrix} = S \begin{pmatrix} V_{xx} \\ V_{xy} \\ V_{yy} \end{pmatrix}$$

When V is steerable these three fields are also steerable. Let us denote $m \in \{0, -2, 2\}$ and:

$$V_m(r, \phi) = \sum_{n=-N}^N W_{m,n}(r) e^{jn\phi}$$

According to the property of the rotation group, the rotation is performed by:

$$\begin{aligned} V_m^\theta(r, \phi) = e^{-jm\theta} V_m(r, \phi - \theta) &= \sum_{n=-N}^N e^{-j(m+n)\theta} W_{m,n}(r) e^{jn\phi} \\ &= \sum_{n=-N}^N e^{-j(m+n)\theta} V_{m,n}(r, \phi) \end{aligned}$$

where the $V_{m,n}(r, \phi)$ are the basis filters. The voting is thus performed in the following manner:

$$\begin{aligned} U_m &= \sum_{n=-N}^N I_{m,n} * V_{m,n} \\ \text{where } I_{m,n}(x, y) &= s(x, y) e^{-j(m+n)\theta(x, y)} \end{aligned}$$

The tensor field is finally reconstructed thanks to the following formula (providing respectively the orientation of the main eigenvector of U , θ_U and the stickness s_U):

$$\begin{aligned}\theta_U &= \frac{1}{2} \text{Arg}(U_{-2}) \\ s_U &= \sqrt{U_{-2}U_2} = |U_{-2}| = |U_2|\end{aligned}$$

Computing U_2 is sufficient because U_{-2} and U_2 are complex conjugates. In [131], the following voting fields are proposed (k being a positive integer, and σ_{TV} a positive scale):

$$V(r, \phi) = \frac{1}{16} e^{-\frac{r^2}{2\sigma_{TV}^2}} \cos^{2k}(\phi) \begin{pmatrix} 1 + \cos(4\phi) & \sin(4\phi) \\ \sin(4\phi) & 1 - \cos(4\phi) \end{pmatrix}$$

The stick vectors of these fields are orthogonal to the classical tensor voting field proposed (cf. [Sec. (2.1.2)]). The voting result is consequently orthogonal to the classical result. The choice $k = 1$ provides the smallest number of basis filters but the broader voting field. This voting field did not provide good performance during our experiments. We chose therefore $k = 2$ as in [131] and we obtained:

$$U_{-2} = (w_0 * \bar{c}_2) + 4(w_2 * c_0) + 6(w_4 * c_2) + 4(w_6 * c_4) + (w_8 * c_6)$$

where the w_i are the basis filters and the c_i contain the input tensor information:

$$\begin{aligned}w_m(r, \phi) &= e^{-\frac{r^2}{2\sigma_{TV}^2}} e^{jm\phi} \\ c_m(x, y) &= s(x, y) e^{-jm\theta(x, y)}\end{aligned}$$

This framework has been used in [97, 131] to enhance catheters detected in fluoroscopic images by an Hessian filter.

2.3 Experimental Validation

This section is devoted to the experimental validation of the two detection approaches presented in the previous section: a method employing classifiers constructed by an asymmetric variant of gentle AdaBoost as detectors and a framework enhancing the response of steerable filters by tensor voting.

We first present the metrics used to quantitatively evaluate the results. Then, the different kinds of fluoroscopic images available in routine clinical context are presented. Finally, the results of our guide-wire detection approaches are presented.

2.3.1 Validation Metrics

In order to evaluate the performance of the proposed methods, we built a database of images where guide-wires were annotated and used as ground truth. We used two metrics to assess the performances of the methods:

- the **Missed Detection Rate** (at d pixels) is the proportion of the ground truth that lies at more than d pixels from the extracted support.
- the **False Detection Rate** (at d pixels) is on the contrary the proportion of the detected support that lies at more than d pixels from the ground truth.

These two metrics (also used in [24]) are respectively related to the recall and the precision. A perfect detection lead to a Missed Detection Rate (MDR) and a False Detection Rate (FDR) equal to 0.

These metrics, however, can only be computed for a given detector threshold. In order to perform a threshold independent comparison of the detectors, we exploited the well-known **Receiver Operating Characteristic** (ROC) curves, that express the True Positive Rate (TPR) in function of the False Positive Rate (FPR) of the considered detectors. The TPR is the ratio of the True Positive (the samples belonging to the positive class that are correctly detected) by the number of positive samples. The FPR is the ratio of the False Positive (the sample belonging to the negative class that are wrongly detected as positive) by the number of negative samples.

2.3.2 Testing Data

The fluoroscopic images currently available in most of the clinical settings are of two types: **raw images** whose intensity ranges from 0 to 2050 and are directly provided by an imaging device (after a simple square-root operation) ; and **8-bits images** that have been processed by a non linear lossy compression algorithm. These algorithms are developed by the fluoroscopic device manufacturers (GE, Siemens, Phillips, Toshiba, ...) in order to reduce the noise and enhance the structures of interest. Because the compression is lossy, one can not relate the intensity of the 8-bits images with the width of the image structures. Moreover, we will call **contrast-enhanced** images the images where the cardiologist has made the blood vessels visible by injecting a radio opaque contrast agent. Last but not least, two kinds of sequences are recorded: **fluoroscopic sequences** are taken with the smallest amount of X-rays, while **record sequences** are obtained with a higher dose. The cardiologist has the legal obligation to record the second type of sequences to keep a trace of the intervention. In particular, he has to record a sequence where the blood vessels have been revealed with contrast agent in order to prove that the patient was suffering from a stenosis. He usually also records a non-injected record sequence at the end of the operation for showing that the stents are correctly deployed.

2.3.3 Guide-wire Detection Results with Boosted Classifier

The first experiment was carried out on a database of *raw* fluoroscopic images (i.e. directly related to the output of the X-rays detector). According to the

Beer-Lambert-Bouguer law, the observed intensity I_p at location p is related to the X-ray intensity send through the patient at location p , I_p^0 by:

$$I_p = I_p^0 e^{-\int_{\mathcal{P}_p} \mu_p(x) dx}$$

where $\mathcal{P}_p$ designates the path of the X-rays through the patient, x the abscissa along this path and $\mu_p(x)$ the absorption at the location corresponding to the abscissa x along $\mathcal{P}_p$. Thus:

$$-\log \frac{I_p}{I_p^0} = \int_{\mathcal{P}_p} \mu_p(x) dx$$

A structure of fixed thickness and absorption (for instance, a guide-wire) induces a constant change of the right term of this formula. Therefore, when a raw image is available, the log is the best way to detect and differentiate structures. For this reason, we exploited features (Difference of Gaussians, steerable filters) computed on the log-image.

On the other hand, the noise (that results from X-ray diffusion and quantum noise and can be modeled by a white noise for fluoroscopic applications, because the number of photons collected at each detector part is sufficient) is modified by this operation. The signal-to-noise of the region of the fluoroscopic images where the measured intensity is the lowest is finally reduced in comparison to the regions of high intensity. For this reason, we also considered features extracted from the raw image.

In [132], the input features of the classifier were only provided by the raw image and by steerable filters for ridge detection [64] of second and forth order, at different scales, and we incorporated responses at 90° and 45° with respect to the orientation providing the best filter response. Then, we also considered the log-image, and we introduced an edge detector, Difference of Gaussians, and variances of the subtraction image (the current frame minus the next one in the fluoroscopic sequence) as well as variance of the difference of Gaussians [133]. In this manner, we tried to discard static confusing structures, to take the background intensity into account and to discard the rib borders.

The quantitative results that we obtained for both approaches are presented by the figure [Fig. (2.4)]. The second approach, as expected, produces the best results. This approach leads however to a high computational burden. Given that the memory and the time required for the detector training increase linearly with the number of features, the results obtained with the second approach can be considered as the best possible with the adopted detectors.

2.3.4 Guide-wire Detection Results with Steerable Filters

The machine learning-based approaches provide the best results but they require a training that make them difficult to use for real clinical applications, because a new training is necessary for each new kind of image encountered.




Figure 2.4: (a) detection performance (on 5.000.000 points sampled from 47 test images) for the first approach (compared with the oriented filter proposed in [52]) (b) detection performance for the second approach (for 5 million points sampled from 45 raw test images; dashed lines: for 1 million points sampled)

In a second time, we have therefore adopted a more robust approach based on the enhancement of steerable filters response by the tensor voting [97]. The performance were of course lower, but this detector can be adapted quickly to deal with new unseen data.

The figure [Fig. (2.5)] presents the results obtained on three images without confusing structures. Our approach, using the second-order steerable filter introduced in [64] exhibiting the best signal-to-noise ratio, is compared with [97] when extracting an Hessian response at a same scale. For these filters, the results are extremely similar. Using a forth-order steerable filter [64] would certainly slightly improves the results, but these filters require much more computing time, whereas the second order filters are computed like the Hessian [64]. For these reasons, we finally used the second order filters. The figure [Fig. (2.6)] illustrates the results obtained on our tracking database.

2.4 Summary

In this chapter, we described the two main approaches that we have developed to address the detection of the guide-wires.

[132, 133] relies on machine learning and steerable filters [63, 64]. The principle of this approach is to train a unique (binary) classifier able to separate the guide-wire and the background. The training relies on the Boosting method [109]. We have addressed the extreme asymmetry of the two classes considered - the guide-wire support is hundred times smaller than the background - by considering an asymmetric variant of Gentle Adaboost [110] similar to [126]. This approach provides promising results but it requires a long training process - and in particular a cross-validation for the asymmetry parameter. Therefore, it cannot be adapted to the extremely broad variety of fluoroscopic sequences that are considered during




Figure 2.5: (i) proportion of pixels extracted by Hessian and steerable filters regularized by TV belonging to the wire (estimated with 3 images by ranking pixels by decreasing filter response). (ii) norm of the steerable filter response and of the TV output for an another fluoroscopic image.




Figure 2.6: (a) 8-bits input fluoroscopic images (b) norm of the TV-regularized steerable filter (the contrast has been modified in order to make the guide-wires more visible)

real clinical interventions. This is the reason why we proposed a second approach.

The second approach relies on only one steerable filter [64], whose response is enhanced by the fast variant [97] of Tensor Voting [85, 86]. We need to emphasize that this method of Tensor Voting has been exploited in [131] for enhancing Hessian filter response in order to segment the ablation catheter used to cure cardiac arrhythmias. Our contribution resides in the new choice of stickness tensor field enhanced, and in the new application of the method.

Both approaches transform the input image into a map whose intensity increases with the likelihood of the presence of the guide-wire. In order to perform guide-wire tracking, we have used this support directly (see chapter 4). Our guide-wire segmentation frameworks, by contrast, uses only the pixels that are the most likely to belong to the wire (see chapter 3).

2.5 Discussions

The learning approach shares some principles with the one presented [134] that uses a classical AdaBoost method to train a dendrite detector relying on features similar to the Haar wavelets. Contrary to this work, however, our features are rotations invariants and we have chosen a Boosting method able to handle unbalanced training sets. Our work exhibits also similarities with [65, 104], because it trains a unique structure detector independent of the structure orientation, what is more robust with respect to image noise that can bias the training when it is performed multiple times for different structure orientations. Both methods exploit moreover steerable filters, that are decomposed into a sum of Gaussian derivatives (that can be efficiently convolved with the images). The main difference resides in the machine learning method employed: the authors of [65, 104] exploit an SVM whereas our method rely on Boosting. We have also incorporated other features, like intensity variance. Last but not least, contrary to our features that are independent from each others, the features exploited in [65, 104] are linked together. While our features provide a set of extrema or ratios, their features provide a d -dimensional function depending on the orientation of the structure (d being the number of features). As a result, their approach learns a rotation-invariant response model of the curvilinear structures of interest while our method intrinsically learns this model. It is however unclear if a method is faster than the other one. Our method requires one optimization per Gaussian feature but this optimization is fast for the low-order derivatives. By contrast, [65] performs only one optimization but that implies to rotate the low-order features as many times as the high-order ones. It is moreover difficult to compare the time required by a Boosted classifier and by a SVM.

This comparison, as well as the detailed study of the cascading classifiers [114, 123], lead us to propose three possible improvements of our detector scheme. First, the introduction of combined features that would take into account the relative orientation of the features. One could for instance consider pairs of features and consider as a supplementary feature the response of a filter for the orientation that

is optimal for the other. The classifier could then model the wires in a similar way to [65] - without estimating the structure orientation, though. Second, the orientation could be determined by a regressor exploiting the orientations of the extracted features. Logistic regression (linear [135], additive [120], etc), boosted regression [112] or SVR [101] could be exploited for this task. Finally, using a boosted cascade would accelerate the computations.

An other simple extension of our method would consist in replacing the asymmetric Boosting update by an Adacost update [128]. This would however lead to a more complex system to set up (since we would have to choose the cost of the samples). A combination of RareBoost [129] and Gentle AdaBoost would be better, since it would not introduce supplementary parameters.

Last but not least, exploiting ranking methods, like [111] (or preferably a more recent method able to deal with unbalanced training sets) would provide more reliable results.

Several approaches similar to our second method have been proposed. For instance, blood vessels centerlines (priorly extracted as middle points between two pixels extracted by the Canny edge detector [66]) are reconstructed in [136] ; and the output of a neuron membranes detector is improved in [89]. These methods could be tested for our application. It would also be interesting to try different stickness tensors. There exist indeed many other choices, such as the difference between the best and the worst response or the difference between the best response and a mean or a median response. It is also possible to take more information provided by the detector into account by performing ball voting in addition to the stick tensor voting. This would incorporate the detection scores computed for the direction orthogonal to the best one. One could also replace tensors with more general functions of finite Fourier transforms (steerable functions) in order to perform a sharper voting and to perform more than two kinds of votes (ball vote and stick vote). Last but not least, one could investigate the training of a regressor that would infer the tensor describing the enhanced structure, or that would infer the tensorial votes from local features.

Chapter 3

Curvilinear Structures Segmentation

3.1 State-of-the-art Methods

In the previous chapter we have presented in detail the literature on curvilinear structure detection.

Depending on the the quality of the detector, either the whole output or only a few pixels - likely to belong to the structures - are exploited. The aim of this chapter is to further exploit the detection results towards complete delineation of guide-wire.

Let us first review the methods that are the most related to our guide-wire segmentation approach. As pointed out in the [Sec. 1.3], many computer vision and medical imaging problems are concerned with curvilinear or tubular structures. If we consider the medical imaging field, then we can cite the following set of problems, that are the most encountered:

- **blood vessels segmentation**

Various image modalities allow for the blood cells to be seen and to perform *angiography*. Vessels can be observed directly - like in the case of eye retinal fundus images ([67, 69, 137, 138]) where a precise automatic segmentation of the blood vessels is of prominent interest for the diagnosis of diseases like the retinopathies induced by diabetes or age-related macular degeneration. For more details on retinal image analysis, see for instance the surveys [8, 9].

Angiography refers to the observation of blood vessels with the help of an X-ray imaging device and a radio-opaque contrast agent (injected into blood vessels in order to make them visible). This kind of imagery is used for the diagnosis of strokes, heart attacks and other cardio-vascular diseases. It is also employed during micro-invasive intervention dedicated to cure such diseases, because precise measures are required and direct observations are not possible during this kind of interventions.

Furthermore, MR angiography (employing MRI) and CT angiography (made possible by contrast agent injection during CT scans) allow 3D and even 3D + Time blood vessel segmentation. In all the cases, the aim is to segment tubular structure trees, where the radius decreases significantly from the root to the leafs.

The problem of blood vessel centerline segmentation was addressed in [139], where a single scale curvilinear segmentation methods [133] was considered in a multiscale setting. Numerous variants of multiscale information fusion have been proposed in the literature. For instance, when a model of the intensity profile of the structure is known, it is possible to weight the filter responses according to a function depending on the scale, such that the weighted response is maximized when the true scale of the structure is present. [140] has determined this weights for a second-order Gaussian derivative profile, [138] for a Gaussian profile, [141] for linear and curved blood vessels. Other examples of handling multiple scales can be found in [142, 143, 144]. More details can be found in a recent survey [7], that presents the literature on 3D blood-vessel segmentation. This review is related to the review [145]

dedicated to MR angiography. [146] compares 13 algorithms for coronary centerlines extractions that span most of the approach developed so far.

- **microtubules/actin filaments segmentation**

Microtubules are part of the cytoskeleton and refer to critical metabolic functions such as molecules distribution. They are not stable: new microtubules continuously appear, elongate, shorten and disappear [147]. Alzheimer's disease is related to a perturbation of these dynamics. Consequently, several methods were introduced aiming to segment and track and thus to count microtubules and quantify their dynamic in a cell observed with Total Internal Reflection Fluorescence Microscopy (TIRFM,[148]). This modality measures the excitation of fluorophores by vanishing waves created by the reflection of an incident wave at the membrane of the cell [10, 149, 150]. Some methods explicitly recover the deformation of the microtubules over time and are related to the guide-wire tracking problem presented in the next chapter.

The filaments formed by the assembly of actin proteins allow the cells to move and deform [151, 11]. They are also observed with TIRFM and their segmentation is similar to the delineation of microtubules.

- **neurite segmentation**

In vitro neuron cultures and histological studies are extremely important for neuroscientists. In this context, observations/measurements refer to 3D images obtained with a confocal or an electronic microscope, where the dendrites of the neurons appear as thin filaments or tubular structures. The segmentation of these thin structures is crucial, because it allows to capture interactions between the neurons and to better understand how our brain functions. This problem has been investigated since the beginning of the 80s [152] (manual segmentation of the neurons is infeasible) and a recent survey [12] details the best methods currently used. This field of research has been less explored than blood vessel segmentation, but it is more related to our topic of interest, because the neurites are thin (with respect to the image resolution) like the guide-wires. It is thus most of the time impossible to segment them by exploiting complex structure intensity profile models like for blood vessels, that are segmented very well with models combining blood vessel edge detectors and blood vessel center detectors.

- **brain and muscle fibers segmentation**

Diffusion Tensor Imaging (DTI) and its extensions are imaging modalities derived from Magnetic Resonance Imaging (MRI) that provide a volumetric map of the diffusion of the water molecules within the observed tissue. It relies on the application of high magnetic field gradients (allowing to measure water diffusion) in different directions. Summarizing this data allows to construct either the covariance matrix of the displacements of the water molecules or higher order diffusion profiles in each voxel of the sample volume. The segmentation of these fibers (in particular) and the development of diffusion imaging (in general) are now extensively studied, to the point of

justifying dedicated workshops [153] and challenges [154] in several first rank medical imaging conferences.

[12] classifies neuron segmentation methods into *tracing methods* and *segmentation methods*. Alternatively, [155, 13] adopts as well *sequential tracing methods* and further elaborates on the segmentation methods through a finer classification: *skeletonization*, *minimal cost path*, *minimum spanning tree*, *active contour based tracing* and special tracing algorithms developed for segmenting Neuromuscular Projection Fibers in 3D confocal image stacks.

In blood vessel literature [7] segmentation methods are grouped into four categories: region growing methods (region growing, fast marching), active contours (snakes and implicit contours), centerline based approaches (tracing methods, minimal cost path) and stochastic methods (particle filters and Markov point processes).

The detection/feature part required for the segmentation of thin curvilinear structures has been detailed in the previous chapter. Curve evolution methods such as the spline-based methods are related to curve tracking methods, that are presented in detail in the next chapter. Therefore, let us briefly review the following groups of methods:

- **tracing methods** consist of deterministic sequential tracing methods and their robust/probabilistic counterparts exploiting Multiple Hypotheses Tracking, Kalman Filters and Particle Filters.
- **active contour methods** that regroup snakes and level sets (that we review only briefly, because snakes are described in the next chapter).
- **region growing and shortest paths methods** such as the Dijkstra's algorithm and methods based on region growing, pixel coding, and fast marching.
- grouping methods, graph based frameworks and linking problems solved by global optimization.

During the section, we will mainly focus on methods that have been applied in the field of medical image analysis, even though remote sensing and contours segmentation will also be evoked. We then review methods that specifically addressed the segmentation of guide-wires. The conclusions presented in the last subsection justify our contributions.

3.1.1 Tracing methods

Tracing methods detect *seed points* that are highly likely to belong to the structure of interest and follow the structure step by step. They are also denominated as *chaining methods*, since they construct chains (sometimes called *tracts*) beginning with a seed point, in which one element is inserted per iteration. *Tractography* variants - in the context of diffusion imaging - are part of this framework. Most of

them rely on extensions of tracing algorithms originally developed in the context of CT or MRI imaging to the case of 3D images of tensors. We refer the reader to [156] for an overview of the current trends and issues in tractography. These methods are sometimes called exploratory algorithms, because the regions of interest of the images are progressively explored.

In this section, we separate the numerous tracing methods employed in the past years into three groups:

- **simple tracing** methods, that were the first developed, that are easy to implement and very fast.
- **multiple hypothesis tracking**, that are more robust thanks to a larger exploration of the hypotheses space and the consideration of multiple likely solutions.
- **probabilistic tracking methods**, that estimate a likelihood of each of the hypothesis and finally determine the most likely structure.

Simple Tracing

In the sequel, we will label as **simple tracing methods** the methods described by [Alg. (3)].

Algorithm 3 Simple Tracing Algorithms

1. Seed points extraction
 2. For each seed point
 3. While stopping criterion not fulfilled
 4. perform local structure model fitting
 5. add the obtained structure portion to the current track
 6. estimate the structure direction
 7. shift in that direction
-

Simplicity is the reason for their application. [157] was initially proposed to segment roads in aerial images, the exploratory algorithm [69] was used to segment blood vessel in retinal images, similarly, [152] was used to extract neurites, guide-wires were segmented in [121, 22] ...

This simple tracing method has been notably improved in following works.

[158, 159] counts among the first instances of simple tracing methods that were amended in the field of medical imaging [69]. The latter was developed to address the segmentation of blood vessels in retinal fundus images, that are modeled with a pair of edges belonging to a finite set of directions detected using dedicated kernels. This idea has been extended in [70] to perform neurite segmentation in 3D confocal images, where neurites were modeled with generalized cylinders. In [160], the convolution kernels have been replaced by a median of convolution

filters (normal to the structure) in order to obtain more robust results with respect to structure discontinuities. [161] presents many similarities with [69], and was applied to neurite segmentation. [162] obtains more robust results by enriching the model with the size of the blood vessel. [163] presents a blood vessel tracing method that represents the centerline with a quadratic B-spline, resulting in a smoother segmentation.

In [164, 165, 166], image moments are computed for estimating the local orientations, center and scale of a local cylinder model assumed to describe 3D blood vessels.

The Steger algorithm [157, 14] reaches subpixel accuracy for known structure intensity profiles. The algorithm determines the direction of the structures with the help of the Hessian, and the possibility of extracting the structure width with simple edge detectors (Gaussian first order derivatives) is described. Steger's algorithm was originally developed for roads segmentation in aerial images and successfully used to a broad range of application, including guide-wire segmentation [21, 167].

The work [168] employs super ellipsoids [169] as local blood vessel models. The parameters of these local models are determined using a MAP formulation. The model requires an estimation of the background and blood vessel intensities, that is performed using robust M-estimators.

This method can be related to the ribbon-snake method [170]. In this sense, it is related to the active-shape based segmentations methods presented in section [Sec. (3.1.2)].

Biplanar Guide-wires tracing

3D fluoroscopic imaging devices have been developed with the aim of providing a more complete information during the angioplasty interventions. These apparatus are made of two orthogonal X-ray tubes and two planar X-ray detectors. They are therefore called *biplanar fluoroscopic systems*. The segmentation of the guide-wire in 3D, based on the pair of images provided by these devices is slightly different than our problem: at one hand, segmenting in 3D is easier because there is no crossing between structures, but on the other hand, the 3D reconstruction from the two available 2D projections is not straightforward.

This challenge has been addressed in [122] where the segmentation is performed into both images. The two traces are then used to reconstruct the 3D guide-wire by taking advantage of the fact that a 3D point can be reconstructed as soon as the intersection between the epipolar line corresponding to one point of one trace and the other trace has been determined. In [22], on the contrary, the tracing is performed in 3D and the guide-wire is locally modeled with a line segment. At each step, the half sphere of possible continuations is projected onto the two images in order to perform a line enhancement. The back projection of the results obtained from the two images and their summation indicates the 3D local line-segment

approximating the best the guide-wire.

Adaptative structure tracing

Simple tracing methods have the advantage of being computationally efficient. On the other hand, a part of the structure can be missed as soon as it does not contain a seed point and its link with the other parts is weak enough to make the tracing stop. Several works attempt to address this limitation by employing *adaptive filtering* methods or in general, frameworks that adapt the stopping criterion to the image context. In [39], the main eigenvalue of the structure tensor [38] is considered as a regularized criterion for determining seed points for the tracing. At each step of the simple tracing, a small window surrounding the current location is considered and the segmentation of the blood vessels that is visible in this window is performed by normalized cut [171]. The criterion used to segment the blood vessel is locally adapted. The structure tensor is then used to determine the structure orientation, and to follow the blood vessel. However, dealing with missing information and recovering gaps remains difficult. These adaptive methods are also exposed to the over-segmentation risk: high tolerance on seed points detection could improve segmentation while introducing many false positive segments as well due to noise. A better approach consists in tracking multiple hypotheses and keeping only the best ones.

Algorithms Exploring Multiple Hypotheses

Instead of tracking only one hypothetical structure, it is more robust to generate multiple hypotheses at each tracing step. This approach inspired several methods.

Multiple hypotheses tracking (MHT), described in [172], maintains multiple structure candidates. At each tracing step, a set of hypotheses is generated given the current state of the retained hypotheses (in [173], for instance, the orientation of the hypotheses are equally distributed ; in [174], they are sampled randomly according to an assignation matrix). The likelihood of these samples (according to the image fit) and an update on the current form of the pdf is performed. Clustering the trajectories represented by the remaining proposal allows to detect branch point and delineate the final network structure (see for instance [173]). The main application of MHT is the simultaneous tracking of multiple targets, like in [175] and in [174]. It is also possible to track curvilinear structures as well. MHT has therefore been applied in [176, 173] to blood vessel segmentation.

Probabilistic Tracking Methods

Under the denomination *probabilistic tracking methods*, we designate all the probabilistic methods considering the current state as a random variable. These methods track the structure from the seed points to their tips or to bifurcations like tracing methods, but determine the most likely structure according to the image.




Figure 3.1: Illustration: a Hidden Markov Model. As in the graphical models literature, arrows represent probabilistic dependencies.

Graphical models have been intensively used to address this task.

The most popular graphical models used in that context are the **hidden Markov models** (HMM). In these models (illustrated in [Fig. (3.1)]) the parameters describing the current state S_t (t is the index of the tracing step) form an unknown vector x_t that is estimated according to a vector of observations y_t . These models satisfy the Markovian property, that states that any future state is statistically independent from the past states knowing the current states. Such a property means that the considered stochastic process state depends only on the current state. As a result, the (hidden) states form a Markov Chain.

When modeling a process with a HMM, users can be interested by: **filtering**, that consists in estimating the likelihood of the current state knowing the past and current observation ; **smoothing**, that consists in estimating the likelihood of a past unknown state knowing the past and current observation ; or determining the **most likely trajectory** in the state space knowing the observations. For blood vessel tracking, filtering is the most important output, because it provides the most likely current state. The blood vessel is finally provided by the most likely trajectory.

Depending on the complexity of the state space and on the kind of state transition, different models and different inference methods are employed. In the sequel, we describe three families of HMM that have been used intensively for segmenting curvilinear structures: **Kalman Filters**, **Particles Filters** and **Discrete HMM**.

Kalman Filters and their Extensions

Kalman filter (KF) is a linear-Gaussian model explaining the distribution of observations in time [177]. The interesting property of KF is that they provide a covariance matrix of the state without supplementary computational burden, which describes the accuracy of the obtained state estimate. Kalman filters have been extensively used to track curvilinear structures. In [178, 179], it has been used to track blood vessels in CT angiograms, that were described with a local cylinder model. The **Kalman Filter snake** [180, 181] is a particular case of parametric active contour whose evolution is predicted by a Kalman Filter. It has been used

to track the axis of neurons in electron microscopy images stacks in [182]. Tractography was also addressed with KF, like in [183]. Because the linear Gaussian model is limited, KF were generalized to deal with non linear transition functions f_t . In the Extended Kalman Filter (EKF) [184, 185], the transition function is linearized at each step in order to apply the same equation as the classical KF. EKF has been used for roads tracking in aerial images in [186].

Despite improvement over conventional Kalman Filter, this method is obviously not able to manage correctly highly non-linear transition functions. Unscented Kalman Filter (UKF) [187, 188] tries to model the mean and the variance change directly, thanks to the Unscented Transform. We refer the reader to [189, 190] for more details. UKF were used many times to delineate curvilinear structure, like in [191] for performing tractography.

Particle Filters

Particle filters [192, 193] extend the Kalman Filters to more general classes of transition functions. Instead of tracking one state, this method tracks multiple *particles* corresponding to likely states of the tracked system, that are weighted according to their relative likelihood.

At each iteration of the tracking procedure, new particle states are sampled according to an *importance function*, that is a probability density function depending on the previous particle trajectories and the previous observation like in the well-known CONDENSATION algorithm [194]. In order to prevent particle degeneracy, one often has to re-sample the particles uniformly. We refer the reader to [192], that is a tutorial on particle filters.

[150, 151] employ particle filtering methods for microtubules tracking, and [195] for tractography. Particle filters were used several times to track blood vessels: for instance in [196] to track coronary arteries, in [197] to segment brain arteries and in [198] (that was improved in [199] by the introduction of new sampling scheme). The works [200, 201] also exploit particle filters to track the medial axis of blood vessels (introducing for instance a novel sampling scheme).

Discrete HMM

When the number of states considered during each tracing step is finite, Dynamic Programming (DP) can be used to build tracks. This algorithm is used to solve problems of the following form:

$$\min_{l_i} \sum_{i=1}^N V_i(l_i) + \sum_{i=1}^{N-1} V_{i,i+1}(l_i, l_{i+1})$$

where $i = 1 \dots N$ are the tracing steps, the l_i are integers distinguishing the admissible displacements at each step and the V_i and $V_{i,i+1}$ are called singleton and pairwise costs/potentials. DP works well with this optimization problem: at

each step t , the variable with the smallest index is discarded by setting:

$$V_{t+1}(l_{t+1}) \leftarrow \min_{l_t} [V_t(l_t) + V_{t,t+1}(l_t, l_{t+1})] + V_{t+1}(l_{t+1})$$

keeping a trace of the successive values of the potentials allows to backtrack the entire solution from the optimal value of the last variable. DP has been used to delineate roads in aerial images in [202] and actin filaments in [203].

When the state transitions are penalized by a distance function, exploiting **generalized distance transforms** [204] speeds up the estimation of the best trajectory [205]. Approximate solvers have also been proposed. The method [206, 207], for instance, performs a primal-dual optimization and requires very few pairwise potential computation compared to the classical DP (for which all the possible pairwise potentials have to be computed) for finding a solution. This method provides a hint of a quality measure of the solution without additional cost: when the duality gap is null, the solution provided is guaranteed to be optimal.

3.1.2 Active Contours

The **active contours** approach consists of positioning a parameterized curve until it fits the features of interest in the image. In a metaphor with physics, the users of an active contour define a criterion to minimize that is called *energy* and that contains at least two terms: an *external energy* depending on a map derived from the image (for instance, an edge map, a ridge map, ...) that drives the model towards its optimal placement in the image plane and an *internal energy* that penalizes non-smooth contours. Numerous provisions of this model were considered like for example the a balloon energy (c.f. [208]) acting on the curve area or a length-preserving term as in [29, 21, 33, 167].

Snakes consist of a parameterized curve (using a set of control points with an interpolation strategy) [209]. Their energy - in the most general case - is minimized using gradient-driven or global optimization if the positions of the control points are discretized. By analogy with physics, the derivatives of the different energies are called *forces* (external forces, internal forces, length preserving force, ...) and they drive the points controlling the contour toward a local optimum.

These methods were considerably used in curvilinear segmentation. In [210, 211], an active model is used to segment blood vessels in 3D MRA: the blood vessel centerlines are represented with cubic B-splines and the wall of the blood vessel with a tensorial product of B-spline. Similarly, a cylinder model is used in [212]. The blood vessels are also described with a curved centerline associated with a varying radius (i.e. generalized cylinders) in [213, 214]. Ribbon-snakes were used in [215] to segment road in aerial images.

External energies, such as the Gradient Vector Flow [47] or regional statistics [216] were considered to address the fact that the minimum determined by the snakes is local. Nevertheless control point-based parameterizations suffer from their inability to cope with topological changes. Topology-adaptive snakes (T-

snakes) [217] deal with snakes merging and splitting, in order to segment branching structures.

Level sets

Implicit active contours, also known as **level sets** were first proposed in [218] to address fluid mechanics problem and re-invented in [219]. In this method, the contour is represented implicitly by the zero level set of an evolving function defined on higher dimension. Level sets have received considerable amount of interest, and extended for instance to simultaneous multiple region segmentation, thanks to the *multiphase level sets* framework [220].

Let us briefly review prior art related to the problem of segmentation of curvilinear structures. We will group them into two groups: the approach proposing new level-set methods dedicated to the segmentation of curvilinear/tubular structures, and the approach incorporating advanced internal forces enforcing thin shapes.

[221] (that was applied again in [222]) describes a level set framework adapted to the segmentation of curves. Contrary to classical level sets, that deal with mathematical objects of codimension equal to 1, these ones were designed to treat structures of codimension equal to 2 (i.e. curves in 3D). This method clearly provides better results than the classical method for lung segmentation. [223] (used for satellite and angiogram segmentation) introduces high-order interactions on the level set, that allow to constrain the shape of the enclosed contour to fit a thin branching structure. In [224], this framework is improved by the exploitation of *phase field models*, that have the advantages of implementation easiness and robustness to the initial conditions.

In [45], a level set is driven by a multiscale gradient vector flow (instead of a classical edge map) to segment blood vessels. This framework was extended to multi-scale Hessian-based flow in [225]. In this work, the vesselness [27] measured in a given voxel is distributed on a cylinder oriented according to the main Hessian eigenvector and scaled according to the Hessian scale. The resulting flow field is used to evolve the level set. Similarly, regional intensity consistency terms [226] are combined with edge alignment terms [227] in [228] into a robust scheme. In order to obtain thin segmented structures, the authors of [229] introduced a shape filtering term, that penalizes the number of balls entirely included into the active contour/surface. Last but not least, [230] introduces a capillarity force directly inspired from physics into their level set, in order to make it evolve more deeply into thin vessels and bronchi.

3.1.3 Region Growing and Shortest Path Methods

In this section, we review the main methods based on region growing and shortest distances computations that have been applied to the segmentation of tubular structures. Shortest path methods compute the shortest path through a weighted graph between a root and a target node. Region growing methods start from a set of initial seeds and expand this region by incorporating pixels according to "fea-

ture” proximity and smoothness criteria. These methods are related because many region growing methods are designed to incorporate pixels of increasing geodesic distance, which can be seen as propagating a front of increasing distance. We present first well-known shortest path methods, and describe then region growing / front propagation methods.

Shortest Path Methods

The algorithm described by Dijkstra [231] is the most popular method for computing shortest paths. It maintains a list of *unvisited nodes* ordered according to a so called *tentative distance*. At the beginning step, all the nodes are unvisited and have an infinite tentative distance. At each step of the algorithm (and while target nodes are unvisited), the unvisited node having the smallest tentative distance is considered as the current node. For each unvisited neighbor of the current node, a travel through the current node is considered and the tentative distance is updated if the length of this travel is shorter.

Dijkstra’s algorithm can only be used when the weights/costs are non negative. Its complexity has been reduced by using a Fibonacci heap for storing unvisited nodes, in [232].

Shortest path/minimal cost algorithms has been widely used for curvilinear structures segmentations.

In [233], for instance, the improved shortest path algorithm [234] is exploited to determine the geodesic length of candidate links between seed points extracted along neurons dendrites. In [235], a heuristic similar to the A^* algorithm (that tries to determines the shortest path without computing as much shortest distances as with Dijkstra’s algorithm) is used to segment the medial axis of coronary arteries in CT scans. In [23], a shortest path on hypergraphs is performed to segment guide-wires in fluoroscopic images. In [236], Dijkstra’s algorithm is used for segmenting neuron axon in electron microscopy stacks: axon parts are segmented in each image slice, and shortest path methods allow the authors to match the axons parts from slice to slice.

Region growing methods [237, 238] are iterative methods to aggregate pixels/voxels to an expanding region initialized from a portion of the structure of interest. These algorithms are greedy, incremental and explore only a small part of the data, like the tracing methods. The simplest idea is to use image intensity as a criterion for stopping the region growth. Such an approach fails to cope with weak edges and images corrupted by noise. In [239], a double object/background competitive region growing is proposed in order to address this problem. Otherwise, a feature image can be exploited, such as in [240] where the image is enhanced with the Vessel Enhancing Diffusion scheme [83]. In this work, the number of pixels added to the growing regions is constrained and consequently false blood vessels bifurcations are pruned. [241] introduced also a scheme that produces smoother segmentations using morphological dilations. [242, 243] are examples of advanced region growing method to blood vessel segmentation. Contrary to tracing, the

region growing methods provide an entire, continuous, structure and not only its centerline. The centerline of such a structure is however easy to obtain by skeletonization/thinning [244] or by medial axis extraction [245]. The 3D homotopy preserving thinning algorithm [246] has thus been used in [242] to provide the skeleton of hepatic vasculature networks. We refer the reader to the review [247] for a detailed insight on skeletonization methods. The more recent review [248] details also a list of properties that a good skeletonization method is expected to exhibit.

Front propagation methods evolve, and gather at each evolution step the traversed locations into the segmented structure. Due to their simplicity and the fact that the smoothness of the front can be controlled, they have been extensively used, for instance for blood vessel segmentation [249]. [250] is a wavefront propagation algorithm applied for blood vessels segmentation. This last work has been made robust with respect to vessel radius variations in [251] by introducing a supplementary dimension corresponding to the blood vessel scale for the propagation. In order to obtain smooth fronts and reject small vessels and false detections, some methods correct the front during the propagation. This is for instance the case of [252] where the evolving surface is topologically constrained.

Fast Marching is a particular case of shortest path front propagation algorithm. It was proposed in [253, 254, 255, 256]. This method determines the solution of the **Eikonal equation**:

$$|\nabla \mathcal{U}(x, y)| = F(x, y) \quad \forall (x, y) \in I$$

where $\mathcal{U}$ - called the *action map* - is the geodesic distance (from a set of roots) that has to be computed, and F is a positive function (both are defined on the image I , and the action map is null at the roots).

Fast Marching was used to segment roads and blood vessels in [257, 258] as well as microtubules [10]. Many approximated algorithms storing nodes in an incomplete order have been proposed, like the algorithm [259] that stores nodes distances in a map structure. It is worth noting that the estimation of the geodesic distance in the regions where the curvature is high can be partially corrected, like for instance proposed in [260]. The computational time can be considerably reduced by adopting a *freezing* principle, as described in [261]. This idea was combined with an anisotropic version of fast marching in [262]. It is worth noting that fast marching has recently been made anisotropic. The work [263] describes a new scheme for performing fast marching on Riemannian manifolds. Last but not least, given the boundaries of a structure are known, fast marching can be used to determine the location of the centerline [264].

Voxel coding [265, 266] bridges the gap between front propagation and fast marching methods. This method evolves a front while marking visited locations with the index of the step corresponding to time of the first visit. It can thus be seen as a fast marching computing a shortest L_1 distance. Due to its simplicity, it has been used frequently like for example in neuron segmentation in [267].

The centerlines of the curvilinear structures provided by front propagation, voxel coding or fast marching are extracted by backtracking [249, 267]. When




Figure 3.2: Illustration: one of the curves considered to compute the Elastica

the centerline is the only output of interest, adding a supplementary dimension corresponding to the radius of the structure in the search space provides better results [268].

3.1.4 Grouping Methods

Early computer vision approaches sought inspiration from the mechanisms of the human vision. The human visual system is able to reconstruct missing non-visible parts of an object, and has the ability to complete/generate/inpaint small missing parts of edges and regions. The aforementioned characteristic is the origin of many famous visual illusions, and was described using the *good continuation criteria* of the Gestalt's psychology [90, 91, 92].

The Euler's Elastica has been proposed in [269] to serve as a good continuation criterion measuring if a point p where the contour has a local derivative τ_p and a point q with derivative τ_q should be linked together according to:

$$E(C) = \int_C \alpha + \beta \kappa^2$$

where α and β are constant parameters and C is a curve linking p and q such that $C'(0) = \tau_p$, $C'(1) = -\tau_q$ and κ is the local curvature [Fig. (3.2)].

We refer the reader to [270] for a review of the possible applications of the Elastica. Several approximations of this criterion have been proposed. For instance, the scale-invariant criterion [271] is the following:

$$e(p, \tau_p, q, \tau_q) = 4 (\theta_p^2 + \theta_q^2 - \theta_p \theta_q)$$

where the angle θ_p and θ_q are the angle between $\vec{pq}$ and τ_p and τ_q as illustrated on [Fig. (3.2)]. This criterion is made scale dependent by multiplication with the Euclidean distance between p and q . It is always greater than zero and smaller than 12π . Though defined for any angle, this criterion is only valid for small angles: the limit when θ_p is fixed and when θ_q tends to π is not the same as the limit when it tends to $-\pi$. In the sequel, however, we use this approximated criterion.

Hierarchical grouping consists in aggregating primitives step by step towards complete recovery of the structure of interest. In [24] the grouping is driven by the output of a boosted classifier [Sec. (3.1.5)]. Hierarchical grouping is greedy (and

thus fast). The same linking problems can be addressed using global optimization methods over appropriately defined graphs towards the optimal linking.

The most promising grouping approaches are graph-based: a graph linking neighboring structure parts is built and the optimal structure is determined by a global optimization performed on this graph. Numerous optimization methods were proposed. The following list is not exhaustive:

Fractional programming [272] is used for extracting the circuit exhibiting the minimal rational cost in a weighted directed graph (the circuit with minimal cost is commonly extracted by iterative applications of the Bell-Ford algorithm). This approach is known as the Ratio Contour in the completion literature [273] and can be used for open contours as well [274]. In [275], the multiple-objects segmentation problem is addressed as a (contour) min-cover problem, that is solved incrementally with fractional programming. More precisely, edgels are extracted, and a set of objects being defined according with these contour portions is searched. The search is addressed by a greedy method adding at each step the curve minimizing the ratio-cost based on the costs of the edgels and on a value counting the number of edgels not covered by the union of objects contours.

Markov Random Fields (MRFs) are undirected graphical models, verifying the Markov independence property: each random variable is conditionally independent of the others knowing the value of its neighbors in the graph associated with the MRF. The optimal solution corresponds to the maximal likelihood. Since the configuration providing the maximal likelihood provides the minimal minus log-likelihood, MRF are commonly expressed into the following optimization problem:

$$\min \sum_i V_i(l_i) + \sum_{(i,j) \in E} V_{i,j}(l_i, l_j) + \sum_{c \in \mathcal{C}} V_c(\{l_i | i \in c\})$$

where l_i denotes the realization of the random variable x_i part of the graphical model, E denotes the set of pairwise relations between variables, $\mathcal{C}$ contains higher order cliques (involving more than two random variables) and the V_p functions are called *potentials*: *singleton potentials* when they involve one variable, *pairwise potentials* when two variables are considered, and *high-order potentials* otherwise. We refer the reader to [276] for a recent review of MRF methods and their applications in image analysis. Roads were extracted by using the MRF framework in [15, 277], and linear primitives are grouped together in [278]. Multiple contours made of linked edgels were also segmented in [279].

Conditional Random Fields (CRFs) [280] are also graphical models, but contrary to MRFs, the random variables *conditionned by the observations* obey a Markovian assumption. CRFs model conditional probabilities whereas MRFs model joint probabilities, and require therefore prior probabilities. CRFs were used for contour completion in [281].

Marked Point process were used in [282] for segmenting roads networks in remote sensing images and later in [283] for segmenting the coronary tree. Such an approach leads classically to apply Monte Carlo simulations, for instance the *Reversible Jump Monte Carlo Markov Chains* (RJMCMC) [284] like in [282]. Other

Jump processes devoted to the segmentation of curves were proposed, like the *Metropolized Gibbs samplers* [285] that was designed to parse an image into a set of regions and a set of mutually occluding curves.

The optimization of graphical models was addressed with a variety of methods: generic solvers are exploited (for fractional programming, MRFs, CRFs) while for hard optimization problems relaxations were considered into a constrained linear or quadratic mixed integer problem, or even into a linear or quadratic integer program. Primal-dual optimization, Lagrangian based resolution and interior point methods are for example used in the later case.

3.1.5 Guide-wire Segmentation Methods

The spread of real-time high resolution fluoroscopic imaging systems in the end of the 90's and in particular of the flat panels in the beginning of 00's motivated the development of guide-wire segmentation methods.

[121] is one of the first one in that domain, where the segmentation is performed by a tracing method based on Hough Transform. Since then, many frameworks spanning roughly all the families of curvilinear structure delineation methods have been proposed. In [21], the Steger algorithm [14] is employed to extract tracts of good quality (subpixel accuracy) to which a B-splines contour approximation are then fitted. [23] employs an hypergraph shortest path method to delineate the guide-wires. In [24], line-segments belonging to the guide-wire are first extracted by a probabilistic boosting tree-based detector. These line-segments are then grouped together, according to the output of a chain of boosted classifiers, that takes lengths, relatives orientations and distances into account to decide when the linking should be done.

The problem of catheter segmentation, and in particular, of the segmentation of the ablation catheters used for curing atrial fibrillation was also investigated [131, 286]. Catheter electrodes are detected using the blobness and grouped, while the catheter body is extracted by thinning a tensor-voting enhanced image in [131]. In [286] the catheter tip is detected using a boosted cascade classifier, and the catheter body is extracted by fast marching.

3.1.6 Conclusion

In this section, we have reviewed the families of methods used for segmenting curvilinear structures and in particular for delineating biological and medical elongated structures such as neurons, dendrites, blood vessels and bronchi. Extracting contours is a major topic in computer vision. For this reason, the literature related to our work is extremely rich.

The challenges related to the segmentation of the guide-wires in noisy fluoroscopic cardiac images dictate our approach. Grouping methods seem to be very promising due to their robustness and due to their good optimality guarantees.

In particular, hierarchical methods are interesting because they allow to separate the low-level concerns (such as determining primitives strongly supported by the image) from the high-level ones (such as recovering blurred/occluded parts of the wire). In [24], however, the hierarchy contains many levels and a training is carried out at each level. Such a complexity is not easy to handle and raises overfitting concerns. For these reasons, we finally decided to base our contribution on a bottom up global optimization method.

3.2 Contributions to Curvilinear Structures Segmentation

The following sections present our contributions to the segmentation of guide-wires. We have adopted a bottom-up approach that takes as input a set of pixels likely to belong to the guide-wire, that extracts primitives from this point cloud and finally delineates the guide-wire by linking the primitives together. These methods take into account the presence of outliers.

Two clustering frameworks are presented. Given a unorganized set of points provided by the guide-wire detectors, we either partition them into elongated groups of pixels or consider the interactions at the interfaces between clusters, which lead us to consider a different framework (MRF with label costs [287, 288]).

The next section describes the linking criteria that we have used during our experiments. We describe finally two methods linking the previous primitives according to the linking criteria, towards segmenting the entire guide-wires. All the components and their alternatives are presented by the figure [Fig. (3.3)]

3.2.1 Most Related Work

Our first primitive linking methods (involving linear programming and supermodular MRF) shares common principles with the MRF and CRF based segmentation approaches [15, 277, 278, 281] and to a certain extend to the approaches based on marked point processes [282]. The recent work [289], that is devoted to the segmentation of trees, relies on a network flow optimization solving a mixed integer quadratic program, that could also perhaps be adapted to the segmentation of wires.

The main difference (and challenge) of the guide-wires segmentation with respect to these most related works resides in the fact that we have to segment curvilinear structures and not a generic network structure, what means that the connectivity of the structure has to be constrained (which lead, as demonstrated in the Appendix A.1, to a NP-hard task). Our last multiple structures segmentation scheme addresses the main limitations of our first one, except that there are still no guarantees on the global optimality of the provided solution.




Figure 3.3: Illustration of the pipeline adopted for the segmentation of guide-wires. Two alternatives are proposed for the primitive extraction and for their linking.

3.2.2 Local Line-Segments Extraction via Clustering

The method presented in this section employs discrete optimization and robust methods to extract line-segments from the input set of pixels: the point cloud is first partitioned into elongated clusters, that are assumed to cover linear parts of the guide-wire and on which line-segments are fitted during a second step.

Linear Programming based Stability Clustering

Clustering has been addressed in many different ways, depending on the kind of elements considered, on the similarity metric considered and on the description of the groups. Most of them cast clustering as a minimization problem like for instance the family of the K-means methods [290], that describe the partition by providing the location of one center per group of elements, assuming that a group gathers all the elements that are nearer to its center than to the centers of the other groups.

The main limitation for most of the clustering methods is that they are unable to determine the optimal number of clusters. It is often addressed by considering different numbers of clusters and keeping the best result, when adding a term penalizing the number of clusters, such as the Akaike Information Criterion (AIC) [291] or the Bayesian Information Criterion (BIC) [292], a strategy that can be time consuming. Furthermore, most of the existing iterative methods exhibit a second important limitation that is their dependency on the initial conditions (number of clusters, cluster centers, etc). It is thus necessary to run them several times with different initializations for increasing the likelihood of getting the optimal result.

In our framework, the guide-wire detector provides a set of points likely to be part of the wire. We need a method able to cluster this set of 2D elements into elongated groups corresponding to linear parts of the guide-wire. The number of clusters is not known in advance, since it is related to the length of the wire.

Clustering based on Linear Programming stabilities [17] addresses the requirements of our framework. It can provide the optimal number of clusters and the assignments of the observations to them while being able to encode arbitrary distance functions between the cluster center and the associated observations. In this formulation a group is described by a center and its members. Let us denote with p_i the points chosen according to the detector output, and with $d_{i,j}$ a measure of distance between p_i and p_j . Let us introduce a list of costs $d_{i,i}$ penalizing the choice of x_i as a cluster center. Let us denote with $x_{i,j}$ a boolean variable equal to 1 when p_j belongs to a cluster of center p_i , and with $x_{i,i}$ a boolean variable equal to 1 if and only if p_i is a cluster center, then [17] aims at minimizing:

$$E_{clustering}(x_{i,j}) = \sum_{i,j} x_{i,j} d_{i,j}$$

$$\text{s.t.} \quad \begin{cases} x_{i,j} \leq x_{j,j} & \forall i, j \\ \sum_j x_{i,j} = 1 & \forall i, j \end{cases}$$

Such a formulation can incorporate any positive function d and eventually varying cluster creation costs depending on the observations. When a local structure orientation is extracted, one can incorporate it into $d_{i,j}$ for producing elongated clusters, whose members orientation is compatible with their location with respect to the cluster center.

The AIC corresponds to the choice:

$$d_{i,i} = 1 \quad \forall i$$

and the BIC, to

$$d_{i,i} = \frac{1}{2} \log(P) \quad \forall i$$

where P is the number of points extracted. [132, 133, 293] considered constant penalties $d_{i,i}$ similar to the previous criteria. The following section presents a better choice that allows to extract clusters at different scales at the same time.

Multi-scale Elongated Clusters Extraction in the Presence of Noise

Two important aspects have to be taken into account when choosing/designing a clustering method able to reliably extract local parts of the guide-wire. Firstly, the fluoroscopic images are noisy and contain numerous confusing structures, like rib borders or staples. As a result, a broad majority of detected pixels are not part of the guide-wire. Secondly, the guide-wires can have slightly different widths from clinical configuration to another, and the diameter of the artery where it is pushed varies significantly.

Let us first consider a set p_i of pixels been extracted by a set of detectors of different scales (σ_i) along with a local wire orientation θ_i . The distance function measuring consistency between two extracted points that we chose combines a geometrical part inversely proportional to the scale of the cluster center and a cluster center penalty proportional to the square of this scale:

$$\begin{cases} d(i, j) &= \frac{d(p_i, p_j)}{\sigma_j} + \alpha e(p_i, \theta_i, p_j, \theta_j) \\ d(j, j) &= \sigma_j^2 K^{MCE} \end{cases}$$

where $d(\cdot, \cdot)$ is the Euclidean distance, α and K^{MCE} are parameters and $e(p_i, \theta_i, p_j, \theta_j)$ is a scale invariant value increasing when the locations of p_i and p_j are not compatible with the local orientation estimations θ_i, θ_j . During our experiments, we chose for $e(\cdot, \cdot, \cdot, \cdot)$ the scale-invariant Elastica criterion [271], that was described in [Sec. (3.1.4)]. With this choices, clusters obtained are σ_j times larger and contain σ_j^2 times more elements when their center has been extracted at scale σ_j , as expected for a multiscale approach. A virtual point gathering outliers, denoted with $\emptyset$, was also introduced and its penalty set to zero. Considering a constant distance $d_\emptyset$ between the point $\emptyset$ and all the p_i forced isolated points at a distance of more than $d_\emptyset$ from any cluster center to be considered as outliers. Since the method [17] is a primal-dual scheme, when it converges to a configuration with a dual gap equal to zero, we have the guarantee that the clusters and outliers, that have been determined are optimal with respect to the design of the objective function.

Relation with the medialness

The medialness [294, 295] relying on the characterization of shapes with their medial axes [296, 297], is related to our clustering. In these frameworks, the centerlines of the structures are indeed the centers of circles of different radii, that are tangent to the border of the structure and the aim is to find the points x belonging to the centerlines, together with the local radius σ_x of the structure. In our method, the centers of the clusters are found together with the optimal scale through the optimization of a criterion over the points chosen as cluster members. One could consider that our method is a discrete, robust, way to extract points belonging to the medial axes. We refer the reader to [245] for more insight on medial representations and their application in computer vision.

As it was done for the medialness, for instance in [141], our framework could also be enriched by the introduction of a model of the structures that are segmented. The more general setting would describe the weighted distances with:

$$d(i, j) = f(v_i, p_j, \theta_j, \sigma_j)$$

where v_i is a feature vector extracted at p_i (containing for instance a structure orientation estimation at p_i , the output of a structure border detector at a scale related to $\sigma_j, \dots$) and f is the minus log likelihood of the probability of p_i being part of a structure with a centerline containing p_j . The only limitation of such a

model is that the interactions between the points of the clusters are not taken into account: when the cluster centers are fixed, a point is assigned to a cluster center independently from the others. This limitation is described in detail in the section 3.2.3, where a more generic framework is examined.

Line-segments Fitting

The extracted partition of points is assumed to correspond to linear parts of the guide-wire. They still contain noisy/outliers pixels, although most of them have been assigned to the outlier class $\emptyset$. In order to fit one line-segment per cluster, while handling partial presence of outliers, we used a variant of the well-known RANSAC algorithm [298], called MSAC in [299], introduced in [18].

Robust model parameters estimation can be performed by at least two different approaches [299]. The first one, employed by the RANSAC algorithm, consists in building models explaining the observed data - each time from a subset of data - and computing a score reflecting the fit of the model to the remaining data. The score takes only into account the data correctly explained by the model (for which the error is under an arbitrarily threshold). The outliers are not taken into account making the corresponding score robust with respect to their presence. More precisely, let us denote with $\mathcal{M}$ a model estimated with the observed data p_i (for instance a set of points in $\mathbb{R}^N$), $p_i^{\mathcal{M}}$ the data predicted by the model, with $d(.,.)$ a distance function defined between data, and with t a threshold, the score $\mathcal{S}$ considered during a RANSAC procedure is the following:

$$\mathcal{S}(\mathcal{M}) = \sum_i s(p_i, \mathcal{M})$$

where

$$s(p_i, \mathcal{M}) = \begin{cases} 1 & \text{if } d(p_i^{\mathcal{M}}, p_i) < t \\ 0 & \text{otherwise} \end{cases}$$

The second way to enforce robustness consists in replacing the classical score/errors function with *robust estimators*, that put less emphasis on the outliers. MSAC algorithm [18], replaces the individual scores by the following ones:

$$s(p_i, \mathcal{M}) = \begin{cases} 1 - \left(\frac{d(p_i^{\mathcal{M}}, p_i)}{t} \right)^2 & \text{if } d(p_i^{\mathcal{M}}, p_i) < t \\ 0 & \text{otherwise} \end{cases}$$

By putting more emphasis on the data that are explained the best from the model, MSAC performs better than RANSAC. As a threshold, we considered the scale of the center of the cluster.

3.2.3 Elongated Structure Parts Extraction

In this section, we describe an alternative method that we have proposed in [293] with the aim of extracting better primitives. This method extracts first elongated clusters by solving an optimization problem that is more generic than the one solved in the previous section 3.2.2. The centerline of these clusters, that correspond to guide-wire or confusing curvilinear structures parts, are then extracted by a discrete optimization based method. They constitute the output of the method (and are described by sequences of pixels).

Clustering through MRF with Label Costs

The main limitation of the clustering method is that points are assigned to cluster centers independently from each other. As a result, two incoherent points can be placed in the same cluster if they are not too incoherent with the cluster center, and the size of the cluster is limited when Euclidean distance is used in the coherence criterion.

In order to address this limitations, we propose [293] to consider the clustering problem as a minimization of a Markov Random Field with label cost energy [287, 288]. Our method determines first a set of points that are likely to belong to the centerline of a structure using the clustering method for example with strong penalty on the number of line segments/clusters. Let us denote with $c(i), i \in \{1, \dots, k-1\}$ these points. Our method aims at labeling the points p_i with labels $l_i \in \{1, \dots, k\}$, where the label k means that they are outliers, and $l_i = j \neq k$ that they belong to the cluster containing $c(j)$. This is achieved by minimizing the energy:

$$E(l_i) = \sum_i V_i(l_i) + \sum_i \sum_{j \in \mathcal{N}_i} V_{i,j}(l_i, l_j) + \sum_{u=1}^{k-1} V(u)$$

where the set $\mathcal{N}_i$ contains the neighbors of p_i (in order to reduce the computational burden); where the singleton potentials insure that $c(i)$ and its neighbors are labeled with i ; where the pairwise potentials tend to group together coherent points and to classify incoherent points in different clusters; and where the label cost reduces the number of retained clusters. In [293], the following potentials were for instance exploited:

$$\begin{aligned} V_i(l_i) &= \begin{cases} \beta & \text{if } l_i = k \\ \min(d_\alpha(i, c(l_i)), \gamma) & \text{otherwise} \end{cases} \\ V_{i,j}(l_i, l_j) &= \begin{cases} \min(\delta^=(d_\alpha(i, j) - \eta), 0) & \text{if } l_i = l_j \neq k \\ \max(\delta^=(\eta - d_\alpha(i, j)), 0) & \text{if } l_i \neq l_j \\ 0 & \text{if } l_i = l_j = k \end{cases} \\ V(u) &= \begin{cases} 0 & \text{if } \sum_i \delta(l_i - u) = 0 \\ K & \text{otherwise} \end{cases} \end{aligned}$$

where $\beta, \gamma, \eta, \delta^=$ and $\delta^\neq$ are constant positive parameters, and $d_\alpha(i, j)$ combines the Euclidean distance between p_i and p_j with α times the scale invariant Elastica criterion 3.1.4 between them. This model constitutes an extension of the model [17] incorporating pairwise interactions between the members of a cluster, and between the members of two different clusters.

Structure Parts Centerline Extraction

The centerline of each structure parts is then extracted by applying the clustering method presented in the previous section 3.2.2 (that can be seen as a discrete medial point extraction method 3.2.2). These points are then linked by solving a binary optimization problem.

More precisely, let us denote with p_a p_b and p_c three points belonging to the centerline of the guide-wire part, and denote with $x_{a,b}$ a boolean variable equal to 1 if and only if p_a and p_b should be linked together. Let assume that we have determined a linking cost $V_{a,b}$ between p_a and p_b and let define:

$$V_{a,b}(x_{a,b}) = x_{a,b} V_{a,b}$$

Similarly, let us denote with $V_{a,b,c}$ a cost related to the linking of p_a p_b and p_c (in that order) penalizing the curvature of the reconstructed structure, for instance, increasing when the angle $\theta_{a,b,c} \in [0, \pi]$ between $\overrightarrow{p_b p_a}$ and $\overrightarrow{p_b p_c}$ diminishes. In the literature, the curvature is more commonly penalized by using the following potentials (since they correspond exactly to the square of the second order derivatives of the curve considered):

$$V_{a,b,c}(x_{a,b}, x_{b,c}) = [x_{a,b} x_{b,c}] \left[(\overrightarrow{p_a p_b} - \overrightarrow{p_b p_c})^2 \right]$$

However, penalizing the curvature of a curve makes it shrink, since the above potentials depend on the Euclidean distance between p_a p_b and p_c . In [293], we used instead the following potentials:

$$\begin{aligned} V_{a,b}(x_{a,b}) &= [x_{a,b}] [d_\alpha(p_a, p_b) - \epsilon] \\ V_{a,b,c}(x_{a,b}, x_{b,c}) &= [x_{a,b} x_{b,c}] [\mu(\pi - \theta_{a,b,c})^2] \end{aligned}$$

where $d_\alpha(., .)$ sums the Euclidean distance between the points and α times the invariant Elastica criterion 3.1.4 and where ϵ and μ are parameters.

In order to reconstruct centerlines with low curvature, we finally solved:

$$\begin{aligned} \min \quad & \sum_{a,b} V_{a,b}(x_{a,b}) + \sum_{a,b,c} V_{a,b,c}(x_{ab}, x_{bc}) \\ \text{s.t.} \quad & \left\{ \begin{array}{l} \sum_b x_{a,b} \leq 2 \\ \forall a \end{array} \right. \end{aligned}$$

The constraints prevent T-junctions. With our choice of potentials, the cost of a T-junction is greater than:

$$C_{min} = 3 \left(\mu \frac{\pi^2}{9} - \epsilon \right)$$

It is thus possible to prevent T-junctions without considering the constraints by increasing μ (and thus the penalization of the curvature) [293]. The resulting problem is a binary Markov Random field that can be addressed by MRF solvers. One has however to keep in mind that the pairwise potentials are supermodular¹, that are the most difficult to deal with. This is the reason why we chose the QPBO solver [19]. Other solver can however address this kind of problem, like the FastPD solver [207, 206] that exploits a primal-dual scheme, and the TRW-S solver [300], that is a sequential improvement of the TRW algorithm [301].

3.2.4 Exploited Linking Criteria

The output of the afore-mentioned process consists of a number of line segments including parts that do not belong the guide-wire as well as missing parts. The next step consists of developing a method doing reasoning on them. To this end, one first has to introduce the notion of distance between two parts/line segments/end-points. We have examined three families of methods to determine this cost.

Elastica Criterion-based Penalizations

Let us denote with a_2 the end-point of a structure a , whose link with a structure b at end-point b_1 is considered. We assume that normalized tangent vectors v_{a_2} and v_{b_1} are known. Most of our experiments were performed using the following families of costs functions (instead of the non-scale invariant Elastica):

$$C(a_2, b_1) = d(a_2, b_1) + \alpha e(a_1, v_{a_1}, b_1, v_{b_1})$$

where α is a constant parameter, d the Euclidean distance and e the scale invariant Elastica criterion [271] presented in 3.1.4. This allowed us to decouple the penalization of the mis-alignment of the structures and the penalization of the distance between their tips. When dealing with linear structures, we considered scale invariant criteria computed between the centers of the line-segments in order to obtain a cost robust with respect to structures crossings.

Detection-output based Criterion

An alternative to Elastica rely on the boosted guide-wire detector [133]. We turned the output into a likelihood using a logistic function, and then we computed the minus log-likelihood of each pixel of being part of the wire. Under independence assumption, the sum of these per-pixel costs along a path between two points is equal to the minus log-likelihood of this path. Applying a fast marching method [253, 255, 254, 256] allows thus to determine the best link between pairs of primitive extremities, according to the boosted detector.

¹in the framework of binary Markov Random Fields, a potential V is supermodular as soon as: $V(0,0) + V(1,1) \geq V(0,1) + V(1,0)$. This is the case here, since $V_{a,b}(0,0) = V_{a,b}(0,1) = V_{a,b}(1,0) = 0$ and $V_{a,b}(1,1) > 0$

More precisely, as pointed out in [302], if a strong classifier $h(x)$ verifies [110]:

$$h(x) = \frac{1}{2} \log \left[\frac{p(x \in \text{wire}|I)}{p(x \notin \text{wire}|I)} \right] = \frac{1}{2} \log \left[\frac{p(x \in \text{wire}|I)}{1 - p(x \in \text{wire}|I)} \right]$$

then, the likelihood of a given point x of being part of the wire is given by:

$$p(x \in \text{wire}|I) = \frac{1}{1 + e^{-2h(x)}} = \frac{e^{h(x)}}{e^{h(x)} + e^{-h(x)}}$$

This yields to the following minus log-likelihood:

$$\begin{aligned} c(x) &= -\log(p(x \in \text{wire}|I)) \\ &= -h(x) + \log(e^{h(x)} + e^{-h(x)}) \end{aligned}$$

Under pixel independence assumption, the most probable links can be determined by any fast marching algorithm, such as [259].

Learned Criterion

The last alternative consists of learning this costs as suggested in [24]. [132] trains a classifier with Boosting in order to obtain a better linking criterion. This classifier combined geometric information about the primitives to link (the previous Elastica criterion, Euclidean distance) and low level information extracted directly from the image. The output of the classifier can be changed into a linking score through the use of a logistic function as earlier described. Such an approach can be related to the *metric learning* that is currently the subject of growing interest for the machine learning community. We emphasize that our training set is strongly asymmetric. In [132], this problem was addressed by using an asymmetric variant of Adaboost. Replacing the classification by a ranking like [111] would constitute an interesting alternative.

3.2.5 Guide-wire Parts Linking

In the previous sections, we have described various manners of extracting primitives from the unorganized point cloud provided by the guide-wire detector. We have also presented different linking criteria to assess their continuation hypothesis. We describe now the methods that we have developed during this thesis to address complete delineation of the guide-wires.

Linking via simple Boolean Programming

In [132], the primitive linking is cast into a simple integer programming problem. More precisely, let us introduce boolean variable $l_{u,v}$ equal to 1 if and only if the primitive end-points u and v (part of different line segments) should be linked together. Let denote with $\mathcal{Q}$ the set of primitive extremities and with $C(u, v)$ a

cost penalizing the link between u and v (in [132], for instance, the output of a classifier was used as linking cost) and ϵ a positive constant. [132] addresses:

$$\begin{aligned} \min \quad & \sum_{u,v \in \mathcal{Q} \times \mathcal{Q}} l_{u,v} C(u, v) + (1 - l_{u,v}) \epsilon \\ \text{s.t.} \quad & \begin{cases} \sum_v l_{u,v} \leq 1 & \forall u \in \mathcal{Q} \\ l_{u,v} \in \{0, 1\} & \forall u, v \in \mathcal{Q} \times \mathcal{Q} \end{cases} \end{aligned}$$

For a small number of line-segments (less than one hundred) this problem can be solved in a reasonable amount of time (a few seconds on a standard office computer) by the state-of-the art integer programs solvers. This method suffers however from severe limitations. The locality of the criterion considered during the linking is the main one: if $C(u, v)$ is larger than ϵ , u and v are never linked together, even though this link would bridge two guide-wire parts. The framework is unable to recover large missing part of the guide-wire: a local high linking cost breaks a structure into different parts, even though we should obviously link them. Besides, the framework do not allow to control the number of detected structures and it is in theory impossible to avoid the creation of loops.

Robust Single Structure Segmentation by Ordering

In order to cope with the above limitations, [133] proposed a method able to handle outlier segments and order the remaining ones for delineating the guide-wire. Let us denote with I the sequence of inliers oriented line-segments that will be the output of the method, let denote with $|I|$ its size, with $I(i)$ the oriented inlier number i and $I(i)_2$ its second end-point and let denote with O the set of remaining outliers primitives. Our method minimizes:

$$E_{linking}(I, O) = \sum_{i=1}^{|I|-1} C(I(i)_2, I(i+1)_1) + \sum_{o \in O} \mathcal{C}(o)$$

where $\mathcal{C}(\cdot)$ increases with the likelihood of the segment of being a part of the guide-wire (see section 3.2.4 for the linking costs $C(\cdot, \cdot)$). During our experiments, we considered singletons costs proportional to the length of the primitives for the sake of simplicity. Since this problem is NP-hard (see proof in appendix A.1), we addressed it with a local search meta-heuristic: we defined for each configuration (I, O) a neighborhood $\mathcal{N}(I, O)$ and, starting from an initial configuration where all the primitives have been greedily linked together, we adopt the configuration with the lowest energy in the neighborhood of the current configuration until reaching a local minimum. In [133], the neighborhood contains all the configurations resulting from a flip of a subsequence of I , a shift of a subsequence of I , a shift followed by a flip, and all the possible subsequence exchange between I and O . These operations are illustrated in figure [Fig. (3.4)].




Figure 3.4: Operations defining the neighborhood: (i) subsequence flip (ii) subsequence shift. The first number is the index of the primitive and the second one indicates their orientation. Gray color indicates the primitives that are concerned by the operation.

Because the neighborhoods are large and because O was also ordered for reducing the number of subsequences exchange between I and O , we observed during our experiments a fast convergence (in less than ten iterations for a hundred segments). The main limitation of this framework is that it delineates only a single structure, and computational time increases dramatically with the number of segments (because $O(n^4)$ operations are performed at each step, if n is the number of primitives). This limitations are partially addressed in the next sections.

Multiple Structures Delineation by Constrained Ordering

First, we introduced the possibility of constraining the extremity primitives by modifying the initialization function implemented and restricting the operations involving the extremities. Then, we extended our method for segmenting multiple structures by considering a list of n_g inliers sequences L_g (instead of a single one) together with the outliers O , by performing subsequences exchange between all the sequences and by considering two new operations: an operation creating a new inlier sequence when it is interesting to do so, and a function removing the inlier sequences that have lost their interest. Denoting with $|I_g|$ the length of the inlier sequence of primitives I_g and with K^{MSS} a parameter, the new energy minimized is the following one:

$$E_{linking}(I_g, O) = \sum_{g=1}^{n_g} \sum_{i=1}^{|I_g|-1} \mathcal{C}(I_g(i), I_g(i+1)) + \sum_{o \in O} \mathcal{C}(o) + n_g K^{MSS}$$

Like for the clustering, the parameter K^{MSS} acts like an information criterion by penalizing the number of inlier sequences considered.

3.3 Experimental Validation

This section synthesizes the experiments presented in [132, 133, 293, 139]. The first experiments deal with the primitives extraction. The next ones validate the primitives linking.

3.3.1 Experiments on Primitive Extraction

We performed two groups of experiments. The first ones were performed on fluoroscopic images in order to establish the performances of the line-segments extraction method. The second ones were performed on fluoroscopic images and TIRFM images of microtubules in order to validate the delineation of elongated parts.

Line-segments Extraction

The experiments presented in [133] were carried out on a database containing 150 images extracted from 15 fluoroscopic sequences acquired during surgical interventions on 13 patients, that we will denote as *database D1* in the sequel. As explained in the last chapter, during this experiment, the guide-wire support was extracted by Boosted detectors trained by cross-validation. Line-segments are extracted and exploited as primitives during this experiment. For this experiment, no virtual class gathering the outlier had been introduced during the clustering. The criterion used to cluster pixels together was the following:

$$d_{i,j} = \alpha d^2(p_i, p_j) + (1 - \alpha)(d_1^2(p_i, p_j) + d_2^2(p_i, p_j))$$

with:

$$\begin{aligned} d_1(p_i, p_j) &= d(p_i, p_j) \min(\tan(\theta_1), t_{max}) \\ d_2(p_i, p_j) &= d(p_i, p_j) \min(\tan(\theta_2), t_{max}) \end{aligned}$$

where θ_1 and θ_2 are the angles between the line passing by p_i and p_j and the ones given by the local structure orientation estimation (provided by one of the steerable filter used for the detection of the contour) and the constant parameter t_{max} prevent a divergence of $d_{i,j}$ when one of the angle tends to $\pm\frac{\pi}{2}$. α was set to 0.3 and the same cost, set to 90000 was used to penalize the choice of a given point p_i as a cluster center: $d_{i,i} = 90000 \quad \forall i$. Detection and clustering results are provided in figure [Fig. (3.5)].

The quantitative results, that are presented in the table [Tab. (3.1)], indicate that our method is able to retain most of the correct detections (since the missed detection rate is only modestly increased) without generating false detections. Trying to reduce the FDR, we adopt a different approach than [24], that try by contrast to minimize the MDR. Our results are promising: in average, half of the detected points truly belong to the guide-wire whereas less than a fifth of the wire is missing. They are sufficient for the linking.

Elongated Structures Extraction

We have first validated the extraction of elongated structures with a database containing 15 TIRFM images from 5 sequences of images where the microtubules of living cells are observed. The first sequence was used to set the parameters and the remaining ones to perform the test. The following parameters were chosen




Figure 3.5: For two images of the database: 1) input image 2) classification scores 3) detected points and clusters 4) detected segments (colors help to differentiate clusters in 3)

stages	detection	line-segments extraction	edgels [24]
MDR	15.4%	19.3%	7.0 %
FDR	53.2%	53.0%	86 %

Table 3.1: missed detections rate (MDR) and false detections rate (FDR) after detection and after line-segments extraction. The last results are reported for the learning approach [24] (on a different database)

[293]: a distance combining the Euclidean distance with 20 times the Elastica was chosen for all the linking. The candidate segment centers were determined with an outlier distance equal to 200 and a cluster penalty to 3000. We fixed the following parameters for the segments extraction: $(\beta, \gamma, \eta, \delta^=, \delta^{\neq}) = (92, 100, 40, 0.3, 10)$ and the label cost $K = 700$. The parameters for the extraction of centerline points where $(25, 70)$ and we chose $(\epsilon, \mu) = (300, 75)$ for their linking.

[Fig. (3.6)] presents segmentation results. They correspond to a MDR of $12.1\% \pm 5.9\%$ and a FDR $13.0\% \pm 4.5\%$ over the 12 test images, both estimated for a distance of 3 pixels, that is approximately equal to the double of the microtubule width, and with respect to a manually annotated ground truth. These results are both qualitatively and quantitatively promising.

We have then used the guide-wires database to validate our segmentation method. We performed a leave-one-out cross-validation to assess the performance of our algorithm: 15 guide-wires detectors were trained using 75 images from 14 sequences, and applied to 3 images of the remaining one. Extracted point cloud were made of the 6000 points from a dense grid of step 4 which were the most likely to belong to guide-wires. We set α to 5 and we extracted the candidate cluster centers with the following parameters $(150, 4000)$. We fixed the following parameters for the segments extraction: $(\beta, \gamma, \eta, \delta^=, \delta^{\neq}) = (80, 100, 20, 0.25, 10)$ and the label cost $K = 800$. The parameters for the extraction of centerline where $(25, 120)$ and we chose $(\epsilon, \mu) = (200, 75)$ for their linking. Some results are provided by the figure [Fig. (3.7)].

The [Fig. (3.8)] presents a quantitative comparison between our two primitive extraction methods. The detection threshold has been chosen separately for each method. The same linking algorithm based on ordering was applied. Although the FDR is clearly reduced, the MDR increases with our second method, because isolated pixels between strongly detected guide-wire parts are considered as background. The extracted primitives are however more reliable for linking, and we obtained better final performances.

3.3.2 Experiments on Primitive Linking

In this section, we present the experiments that were performed for validating our primitive linking methods.

Linking via simple Boolean Programming

The database $D1$ was used to validate our boolean linking algorithm, and a classifier detecting only the part of the guide-wire that lies outside the catheters was trained. The detection was performed by a boosted classifier exploiting the ridge detectors [64] of order 2, $\mu = 0.0$ scale 1.0 and order 4, $\mu = 0.25$ and scales 0.8, 1.6, 3.0. We estimated the guide-wire orientation with the filter at scale 1.6. Three responses were considered for each filter: the best one, and the ones obtained at 45 and 90 of the best orientation. The best filter [52] was the filter of length 30 and width 3.




Figure 3.6: (i) TIRFM images (ii) final detected microtubules (iii) elongated groups of pixels extracted (with different colors. The red color accounts for the background pixels). The exploitation of local orientation estimation played a crucial role during the partitioning of these point clouds.




Figure 3.7: (1) extracted point cloud (6000 points) (2) extracted segments (distinguished by colors; the red one contains the outliers) (3) parts of the guide-wire finally reconstructed. Centerlines points are marked in white and links in red.

stages	MDR[293]	FDR[293]	MDR[133]	FDR[133]
point cloud	8.5 %	65.7 %	15.4 %	53.2 %
salient clusters	24.3 %	41.5%	-	-
extracted curves	27.7 %	31.4 %	19.3 %	53.0 %
final guide-wire	22.2 %	10.4 %	23.8 %	14.3 %

Figure 3.8: Quantitative comparison between our line-segment based approach [133] and our second approach [293].

The function $C(u, v)$ was the outcome of a classifier that was also trained. The feature space of this classifier was a combination of geometric and photometric features. In terms of geometry, we considered features derived from the euclidean distance between u and v , the corresponding segments' centers, the square of these two distances and min and max lengths of the segments. High level geometric features were considered as well, like the Elastica defined in [271], computed between extremities or between segments centers, and multiplied or not by the corresponding distances and square distances - in order to introduce scale-invariant and scale-dependent curvature constraints. The photometric features are also considered to re introduce low-level information in the process: min and max mean intensity along the segments and terms insuring likely links: the means of the output of the detection along cubic splines linking extremities or centers of the segments (and tangents to segments) - multiplied or not by corresponding distances and square distances to take the length of the cubic curve into account or not.

The figure [Fig. (3.9)] presents some detection, clustering and line-segments extraction results. The figure [Fig. (3.10)] presents quantitative results on the detection performances. The results presented in the table [Tab. (3.11)] indicate that the linking procedure is neither able to significantly reduce the amount of false detections (detecting thus multiple confusing structures) nor able to reduce the amount of missed detection, which means that the detection gaps are not filled. The figure [Fig. (3.12)] presents several final results. The results are characterized by inability of the method to reduce the missed detection rate. A plausible solution is to increase the sensitivity of the detector and thus the false detection rate ; and the performances of the detector, that rejects guide-wire pixels when the background intensity is too low.

Linking by Ordering

We report in this section the results obtained on the database $D1$ with our second primitives linking method, that is based on the search of subsequences orderings. These results were obtained by cross-validation: the classifier was trained and the parameters were set with 75 images from 14 sequences and the performances were obtained on 3 images from the remaining sequence. This experiment was performed for each sequence and the averaged results are presented. The figure [Fig. (3.13)] presents the costs that are computed by our method by performing a fast marching on the classifier output in order to determine likely continuations.




Figure 3.9: Detection, clustering and line-segment fitting results. The rows respectively present: (i) the input images (ii) the output of the classifier (iii) the clustering results and (iv) the corresponding line-segments




Figure 3.10: Detection performances (on 5.000.000 points)

stages	MDR	FDR
detection	25.78 %	83.52 %
segments	27.02 %	76.54 %
linked segments	25.91 %	64.68 %
splines	27.66 %	52.36 %

Figure 3.11: Quantitative performances obtained during the simple linking experiment, that was carried out on 47 images of the *D1* database (MDR : *missed detections rate*, FDR : *false detections rate*)




Figure 3.12: Simple linking results




Figure 3.13: Minimal cumulated costs determined by the approximated fast marching. Colors varies according to the primitive extremity that is used as a seed for the front propagation and according to the distance to this extremity (the brighter the color the higher the cost)

stages	MDR	FDR
detection	15.4%	53.2%
segments extraction	19.3%	53.0%
segments linking	23.8%	14.3%
splines fitting	25.8%	13.2%

Table 3.2: missed detections rate (MDR) and false detections rate (FDR) after each step. The state of the art [24] reports $22 \pm 3\%$ MD and $10 \pm 3\%$ FD (on a different database).

These costs are combined with an Elastica criterion in order to define a linking criterion [133]. The final quantitative results obtained with the parameters [133] are provided by the table [Tab. (3.2)], and several segmentations are showed in the figure [Fig. (3.14)].

The significant improvement with respect to the results provided by the simple linking have two origins: first, the better performances of the detectors 3.3.1 allow to reduce the false detection rate of the detection without missing too large guide-wire parts, and then the ordered linking is able to jump over significant gaps. The linking criterion plays also a positive role.

3.4 Summary

In this chapter, a bottom-up pipeline for the segmentation of the guide-wire has been presented. The input of this pipeline is a set of points selected according to the response of a guide-wire detector. These points are grouped into clusters corresponding to mid-level primitives that are delineated and finally linked together in order to reconstruct the guide-wires.

In section 3.2.2, a graph partitioning algorithm is used to extract dense groups of points, that are approximated with line-segments by a variant of RANSAC. A more generic clustering algorithm is exploited in section 3.2.3, that extracts elongated groups assumed to correspond to long parts of the structures. The centerlines of




Figure 3.14: Final results provided by the ordering method (single structure delineation).

these parts are then extracted and considered as primitives. Then, in section 3.2.4, three different linking criteria have been proposed. The first one is inspired from the perceptual grouping literature. The second exploits the output of the guide-wire detectors (when the guide-wires are detected with boosted classifiers). The last one is itself based on a classifier and is inspired from the hierarchical grouping literature. The delineation of the guide-wires is then studied. After a simple integer programming based approach that suffers from limitations, an ordering based method is presented and then extended to be able to segment multiple guide-wires simultaneously. Aiming to further improve the linking, a method that incorporates into our pipeline the information provided by the segmentation of the blood vessel centerlines in a contrast-enhanced image related to the fluoroscopic image was presented.

3.5 Discussion

3.5.1 Limitations

Both primitive extraction schemes have proven to be effective. The first one is however much simpler and the performance increase provided by the second scheme is not sufficient to justify its further exploitation given the associated computational burden. Besides, our first primitive linking method is not robust enough and the centerlines provided by our supermodular MRF based method have low curvature (due to the high curvature penalization required to avoid the apparition of junctions). Our last segmentation algorithm, in spite of the experimental promising results obtained, does not provide guarantees on the quality of the local energy minimum that it provides.

As a result, we think that line-segment primitives are sufficient for the task of guide-wire segmentation (there is no need to extract curved centerlines, but line-segment length should range in a large interval). Otherwise, it would have been interesting to compare a MRF with third-order cliques with our supermodular MRF. Such a MRF can be either decomposed into a classical MRF following [303] (that would probably have been supermodular) or solved by dual decomposition [304].

3.6 Future Work

One could try to further improve the quality of the extracted primitives, for instance, by exploiting other or new graph-based methods (such as spectral clustering, SVM based clustering ...). Exploiting better curvilinear structure detectors and enhancement scheme (such as Gradient Vector Flow [47], shock filters [80, 81] ...) would be of crucial interest because the poor visual support remain the main challenge.

Other meta-heuristics could also be tested for the delineation of the structures. One should try other well-known solvers used to address the Degree Constrained

Minimal Spanning Tree (see Appendix A.1). A slightly different formulation of this task could probably be related to another graph problem.

A more promising direction would consist in fusing the different components of the pipeline. Reducing the number of components would indeed reduce the number of parameters to setup, as well as increase the robustness of the method. In the ideal case where all the pipeline could be condensed in only one optimization step and global optimality criteria would be guaranteed, the outliers would be rejected in an optimal way and one would perform highly robust segmentations.

Chapter 4

Registration for Guide-Wires Segmentation and Tracking

4.1 Introduction

This chapter deals with geometric and iconic registration methods in the context of guide-wire segmentation and tracking in fluoroscopic images. In particular, we show that geometric methods can be exploited to fuse the information providing by different kinds of fluoroscopic images in order to improve the segmentation of guide-wires. We also show that their tracking can be successfully addressed by combining geometric and iconic methods. Similarly to the tracking of the coronary tree, this application would help the cardiologist during the angioplasty interventions. Tracking of guide-wires has been addressed by a limited amount of methods, but it is related to many other computer vision and medical imaging fields of research. To our knowledge, the most relevant are the following ones.

In computer vision, the **registration** task consists in determining the transformation that registers a source to a target image, according to a similarity metric and a regularization constraint that penalize a class of unauthorized transforms. The rigid registrations are for instance performed by exploiting criteria that penalize all the non-isometric deformations without penalizing the rigid ones. Pose-invariant registration, that is of particular interest for real clinical application, can be achieved with criteria that penalize all the transformation except the affine transforms. Without motion model, the tracking of a deforming object simply consists in registering it from one frame to the next.

Besides, many methods track objects by their contours. This can be achieved with the help of an active contour: a snake when the contour is parametrized (and does not undergo topology changes) [209], or a level-set otherwise [219]. More robustness can be introduced by modeling the displacement of the contour, for instance by performing linear predictions with a Kalman filter [177], as in the Kalman snakes framework [180, 181], or a non linear prediction with the help of a particle filter [194], similarly to [305].

Blood vessels tracking/surgical devices tracking is a critical step in many clinical interventions. Cardiac applications are the most challenging because the cardiac motion is large. It has been achieved either by dense methods (for instance, based on the Lucas-Kanade optical flow [306] as in [307] or based on the image intensity as in [308]) or by geometric methods based on landmarks matching as in [309]. Except guide-wires, many thin surgical devices are tracked by the mean of imaging systems, such as catheters [20], needles used during abdominal surgery [32] and ablation catheters [131, 286].

Microtubules and actin filaments tracking exploit the segmentation of the microtubules and actin filaments presented in the last chapter as an initialization [310, 150]. These methods aims to quantify the elongations of the microtubules and to describe the deformations of the actin filaments.

The tracking of guide-wires is particularly challenging because of the low signal-to-noise ratio and because of the large non-rigid deformations due to cardiac and respiratory motion.

The remainder of this section is organized as follows. We first review the meth-

ods that are currently used for geometric tracking and matching/landmark based tracking. We propose then a matching method that incorporates the information provided by the segmentation of blood vessels in a contrast-enhanced image into our guide-wires segmentation pipeline via graph matching. An iconic-geometric hybrid method for guide-wire tracking is then presented. These two methods are validated and discussions conclude the section.

4.2 State-of-the-art Methods

In this section, we review the main registration methods. In the previous chapter, tracking methods have already been classified according to the number of hypothesis kept in memory during the tracking: the Kalman Filters [177] - that maintain only one hypothesis - have been distinguished from the Multiple Hypotheses Tracking [172] and the Particle Filters [192] - that maintain several hypotheses. We have also classified them according to the kind of tracking model that they exploit: linear-Gaussian for KF, more complex for the generic HMM, that can be addressed by Unscented Kalman Filter [188] and Particle Filters [192]. Therefore, we will focus in this section on the tracking methods that do not incorporate a motion model.

These methods are closely related to the registration methods. They are also closely related to the discrete counterpart of the registration methods: the matching methods, whose purpose is to determine the best match between two sets of primitives/landmarks. In the sequel, we describe these two categories separately: we describe first the **geometric registration methods** that extract salient points (called landmarks) along the structures and track them and we describe then the **iconic registration methods** that determine the deformations by optimizing criteria derived from a dense description of the object to track. All these methods can be made more robust by exploiting a motion model. Such strategy has been employed in all the fields of research: Microtubules were tracked with HMM in [311], actin filaments with particle filters [151]...

4.2.1 Geometric Tracking Methods

Geometric registration methods are based on the tracking of reliably discriminative parts of the object undergoing deformations. When the tracked locations are points, they are often denoted **landmarks** or interest points. Object deformation is then recovered by compiling the trajectories of the landmarks. In this manner, it is possible to exactly determine the motion of a rigid objects and objects made of rigid parts with very few landmarks. Since the only way to track non-rigid objects consists in interpolating the displacements determined for the landmarks, one can only track objects made of elastic parts with this method.

Many different kinds of landmarks can be used. The firsts and most famous ones are the corners, such as provided by the Harris detector [38] and the Shi-

Tomasi (sometimes called Tomasi-Kanade) detector [312, 313]. Both are based on the structure tensor.

In order to increase the reliability of the extracted landmarks, scale space methods were considered. The detector being associated with the SIFT [36, 37] are thus localized as extrema of the Difference of Gaussians over space and scale (the DoG been used to approximate the Laplacian of Gaussian). Affine-invariant detectors were also proposed, with the aim of obtaining landmarks that can be matched in spite of view-point changes. We refer the reader to [314, 315, 316], that compare respectively different interest points detectors, different corners detectors and different affine-invariant regions detectors.

It is possible to group the geometric methods exploiting these landmarks into four main families.

The **Procrustes analysis** can be used when the correspondences between the source and target landmarks are known, which is the case, for instance, when SIFT descriptors are used. In order to obtain reliable correspondences, many variant of the original SIFT descriptor were proposed, such as in [317].

More advanced methods recover the deformations by transforming the landmarks sets and aligning them, such as in the recent work [318]. An important approach consists in considering the points sets as distributions and matching these distributions. Initiated by [319], this seminal idea led to many developments, such as [320, 321], that rely on kernel density correlations. Gaussian Mixtures Models were also intensively used to describe the distributions.

The **Iterative Closest Point** [322, 323] performs iteratively two steps: a step during which the closest points on the target object related to the landmarks extracted from the source object are matched, and a step during which these correspondences are exploited for estimating a transformation between the source and the target, that is applied to the source. ICP has been extended and improved many times (see [324] for a review). For instance, [325] proposed the TPS-RPM algorithms, that iteratively estimates the correspondences with the softassign method and computes the transformation described within a Thin Plate Spline framework.

The **graph matching** methods build graphs from the set of landmarks and match them by minimizing a global energy criterion. Depending on the criterion, different classes of transformations are penalized. For instance, penalizing the change of the inter-landmarks Euclidean distances lead to penalize the non-rigid transformations, and only them. Graph matching has been addressed in many ways: with spectral methods in [326], with quadratic programming in [327], with Markov Random Fields solvers in [328] or with tensor optimization, as in [329, 330]. These two last approaches have the following advantage: they can be extended, by modifying the exploited solver, to deal with higher-order potentials. More insight about the geometric registration methods can be found in [331].

4.2.2 Iconic Tracking Methods

Contrary to geometric, iconic tracking relies on dense displacements fields: this approach determines the displacement of each point of the moving objects, and is therefore based on criteria that are integrated on the whole object surface/volume.

Active Contours

In this section, we briefly review active shapes methods that can be used to track curvilinear and tubular structures. Because level sets undergo topology changes (this property makes them attractive for object segmentation) they are less interesting to the segmentation and the tracking tasks considered in this thesis. We will therefore not detail their description given in the last chapter and focus only on snakes [209].

As explained in the section 3.1.2, a snake is related to an energy made of at least two terms: an external energy that measures if the snake location in the image is prominent (for instance, by integrating the output of a ridge detector) and an internal energy that ensures the smoothness of the contour (for instance, by penalizing the curvature). In the same manner, the speed of the deformation of the contour can be penalized with an appropriate energy. One can also incorporate an a priori knowledge about the contour shape by choosing a particular parametrization. For instance, one can parametrize the contour with a basis provided by a Principal Component Analysis of contours already segmented. The set of parameters leading to the smallest total energy describes the optimal location of the snake in the image. By analogy with physics, the derivatives of the energies with respect to the time are called *forces*.

In order to address tubular structures segmentation and tracking, the generic snakes have been amended with a local width or into generalized cylinders. The simple (open) snake [209] has been extended to incorporate a width information in [170]. The resulting parametric model is called **ribbon-snake**. Such a model is particularly attractive for segmenting roads in aerial images [332], and was therefore applied again in [215, 305]. The same idea is used in [268], where the local width of a curvilinear tubular structure is extracted at the same time as its center-line by performing a 4D fast-marching. With this approach, the authors extracted in a robust manner the coronary trees observed during classical fluoroscopic angiography diagnosis interventions. The generalized cylinders [333], used for instance in [213], rely on the same principle extended for 3D tubular structures segmentation. [7, 334] review many cylinder models that have been proposed.

Snakes are broadly used. Most of the time, the active model is parametrized with cubic B-spline [29, 21, 167]. The 3D equivalent of these models allow to describe the registration of a 3D volume with the displacement of a set of control points, as proposed in the Free Form Deformation (FFD) model [335]. At the beginning, the FFD were not necessarily performed with B-splines [336], but this model gained popularity when splines were introduced [335]. Snakes were recently used for blood vessel segmentation in [213, 214, 46]. In these works, a general-

ized cylinder model is used to model the blood vessels. The segmentation is often performed by a gradient descent on an energy that either enforces intensity homogeneity [213] - like in the Chan and Vese method [226] - or leads to a maximization of a flux through the snake boundary - for instance a Gradient Vector Flow [47] or an Optimally Oriented Flux [119]. The method [337], that was applied to the tracking of coronary arteries, incorporates several of these previous approaches: a multi-resolution displacement field is computed and used as an initialization for a snake-optimization with an external energy based on a GVF.

The methods employed for minimizing the snake energies can be split into two major groups: continuous and discrete methods. We discuss these optimization strategies in detail in the following sections.

Snake Energy Minimization

Consists in performing a gradient descent on the parameters of the snake. Any kind of gradient descent can be used (simple gradient descent, conjugate gradient, Powell's conjugate direction method, Newton methods, Levenberg-Marquardt, ... as well as their stochastic variants). The main limitation of such a strategy is that it provides only a local optimum, and it therefore depends on the initialization. Different approaches were proposed to address this problem. The so-called Ziplock-splines [338] divides thus the set of control points of the contour into one passive set and two active sets. During the optimization, the deformation are propagated from the (actives) extremities of the splines toward the center of the spline, the passive part being at each time optimized with respect to purely geometric constraints. Thus, if the endpoints are correct, the optimization leads to a correct contour. Smoothed external energies have also been proposed, like the Gradient Vector Flow [47], that diffuses edge detection across the image. In spite of these improvements, the presence of elongated confusing structures near to the initialization of the snake can lead to false detections and erroneous tracking.

The discrete methods have been investigated in order to tackle these limitations. These approaches consist in discretizing the admissible displacement of the control points and in using a discrete solver to determine the optimal solution. In particular, when all the energies can be decomposed into a sum of terms that depend at most on two successive control points, the snake optimization can be expressed into a chain-MRF. It is then possible to use a dynamic programming solver to determine the contour displacement [20, 339, 340]. In order to accelerate the computation, the most recent works in that domain [21, 33, 167] exploit sophisticated solvers that converge quickly to a good suboptimal solution without requiring the computation of all the possible values of the terms involving a pair of control points displacements [206, 207, 21, 167] - which considerably reduces the computational burden with respect to the original dynamic programming. If the actual displacement of the control point is part of the ones that were discretized, this approach is much more robust than the continuous one. When the displacements are large, however, quantizing the search space requires too many discrete labels and these methods become less interesting from computational view point.

Coarse-to-fine strategies could potentially address this limitation while reducing complexity and ensuring a good sub-optimal solution.

Iconic Guide-wire Tracking Methods

Most of the guide-wire tracking methods that has been applied so far rely on iconic registration [341, 26, 29, 31, 21, 167].

In [341, 26, 29], the guide-wire is represented with a B-spline, and tracked into the snake-spline framework. The main differences of the two methods reside in the external force, that is based on Hessian eigenvectors and eigenvalues in [341, 26] and on phase congruency [28] in [29]. The smoothness term is also different since [29] introduces a length preserving term. Last but not least, the optimization is addressed with a different scheme: the energy minimization is performed by the Powell's conjugate direction method in [341, 26], whereas it is performed by gradient descent in [29].

In [31], the tracking is expressed into a Bayesian framework and a coarse-to-fine procedure is employed. The guide-wire is split into three parts: catheter tip, guide-wire body and guide-wire tip, and the tracking is made more robust by combining, for each part, the output of a trained classifier with an appearance models.

In [21, 167], the optimal displacement of the control point is obtained by solving a Markov Random Field. The same framework has been endowed with third order potentials in order to encode high level rigid transformation parameters during needle tracking [32]. It has been reinforced by a motion model in [33].

[26] was adapted [122] in the context of tracking guide-wires observed with bi planar fluoroscopic devices in the following manner: both 2D guide-wire *traces* (i.e. projections onto a planar X-ray detector) were tracked with the method [26] and the wire was reconstructed in 3D by exploiting the geometry of the device (that was assumed to be calibrated) and stereo-vision concepts, such as epipolar constraints. More precisely: knowing the projection matrices, [122] reconstructed each 3D point of the wire from a pair of 2D points that were part of the wire traces. Choosing the first point in the first projection, they determined the epipolar line in the second image where the second point of the pair is obliged to lie. The intersection of the guide-wire trace in the second image with this line provided them the second point of the pair, allowing them to obtain the 3D coordinate of the guide-wire point.

4.3 Linking Improvement via Graph Matching

As highlighted previously, in clinical settings the combination of the motion blur and the image noise makes large parts of the guide-wire invisible.

Fortunately, fluoroscopic sequences are taken under different conditions during a standard clinical intervention. The cardiologist records a fluoroscopic sequence where the coronary arteries have been enhanced by a contrast agent before (and after) the stent deployment. In these sequences, the guide-wire is not visible due

to the contrast agent but the arteries are easy to segment. Given that the guide-wire lies in an artery, a segmentation of the contrast-enhanced sequences can be employed, either for enhancing the parts of the wire that are hard to see or for improving the links between the visible parts. In this work, we focus on the second approach. A set of line-segments approximatively covering the centerlines of the blood vessels is extracted from a contrast enhanced image and matched to the line segment extracted from the fluoroscopic image. Considering paths between neighboring line-segments of the fluoroscopic image through the contrast-enhanced image when they provide a smaller distance, we aim at bridging missing parts of the guide-wire [139].

The main challenge encountered when performing blood vessels segmentation originates in their diameters variation (from more than 2-3 millimeters for the main arteries to only 0.4-0.6 millimeters for the thinnest visible ones). This challenge was addressed by our multi-scale scheme 3.2.2.

The purpose of our graph matching is to simultaneously determine as much interesting shortcuts as possible between neighboring line-segments of the fluoroscopic image through the contrast-enhanced image. In [139], we describe a framework where the sum of the length of all the shortcuts resulting from the matching is minimized. Let us denote with $a_i = (a_i^1, a_i^2)$ a line-segment extracted in the fluoroscopic image, and with $b_j = (b_j^1, b_j^2)$ a centerline part extracted in the contrast-enhanced image. A graph G^A containing one node related to each oriented line-segment a_i is constructed. The edges of G^A are weighted by the costs of the links between the line-segments corresponding to its nodes. Let us denote with $A_{i,j}$ the matrix storing these weights. A graph G^B is built for the b_j in a similar manner. The matrix of the shortest paths of G^B , $C_{i,j}$ is computed with the Floyd-Warshall algorithm (we assume here that there are no negative circuits in G^B thanks to a sufficient penalization of the curvature of the links). Let us define the following travel costs between the two graphs: $V_\phi(i, j)$ from G^A to G^B and $V_\psi(j, i)$ from G^B to G^A (for $i \in G^A$ and $j \in G^B$). The matching from G^A to G^B is described by the set of boolean variables $\phi_{i,j}$ indicating if $i \in G^A$ should be linked with $j \in G^B$. Similarly, $\psi_{k,l}$ is a boolean variable indicating if $k \in G^B$ should be mapped to $l \in G^A$.

The cost of the most interesting shortcut between i and l in G^A through G^B starting from i , arriving in G^B at j , leaving G^B at k and rejoining l - illustrated on figure [Fig. (4.1)] - is given by:

$$V_\phi(i, j) + C_{j,k} + V_\psi(k, l)$$

With our notations, minimizing the sum of the shortcuts under the constraints that a point in G^A is mapped to exactly one point in G^B and that a point in G^B




Figure 4.1: (i) line-segments extracted from one of the fluoroscopic image of the sequence and the line-segments assumed to delineate the blood vessel centerline of a related contrast-injected image (ii) illustration of the shortcuts considered by the graph matching

is mapped to exactly one point in G^A leads to the following boolean program:

$$\begin{aligned} & \min_{\phi_{i,j}, \psi_{k,l}} \sum_{i,l \in G^A} \sum_{j,k \in G^B} \phi_{i,j} [V_\phi(i, j) + C_{j,k} + V_\psi(k, l)] \psi_{k,l} \\ \text{s.t. } & \begin{cases} \sum_{j \in G^B} \phi_{i,j} = 1 & \forall i \in G^A \\ \sum_{l \in G^A} \psi_{k,l} = 1 & \forall k \in G^B \\ i \text{ and } l \text{ are neighbors in } G^A \\ \phi_{i,j} \in \{0, 1\} & \forall i \in G^A \quad \forall j \in G^B \\ \psi_{k,l} \in \{0, 1\} & \forall k \in G^B \quad \forall l \in G^A \end{cases} \end{aligned}$$

Since we want to recover missing guide-wire parts, we will cope with missing correspondences during the matching. A simple way to address this problem consists in introducing variables $\phi_{i,\emptyset} \forall i \in G^A$ and $\psi_{j,\emptyset} \forall j \in G^B$ indicating that i (respectively j) are not matched.

In order to simplify the formulas, let us denote:

$$\begin{aligned} S &= G^A \times G^B \cup G^A \times \{\emptyset\} \cup G^B \times G^A \cup G^B \times \{\emptyset\} \\ S_i &= \{i\} \times G^B \cup \{(i, \emptyset)\} & \forall i \in G^A \\ S^j &= \{j\} \times G^A \cup \{(j, \emptyset)\} & \forall j \in G^B \end{aligned}$$

The S_i and S^j partition S . Let us also introduce:

$$M_{i,j}^{k,l} = [V_\phi(i, j) + C_{j,k} + V_\psi(k, l)]$$

Let us now relax the constraints on the $\phi_{i,j}$ and $\psi_{k,l}$ (denoting with $\tilde{\phi}_{i,j}$ and $\tilde{\psi}_{k,l}$ the real variables associated with the $\phi_{i,j}$ and $\psi_{k,l}$). The constraints are now:

$$\begin{aligned} \sum_{j \in G^B \cup \{\emptyset\}} \tilde{\phi}_{i,j} &= 1 & \forall i \in G^A \\ \sum_{l \in G^A \cup \{\emptyset\}} \tilde{\psi}_{k,l} &= 1 & \forall k \in G^B \end{aligned}$$

Let us store the matching variables $\tilde{\phi}_{i,j}$ and $\tilde{\psi}_{k,l}$ into a vector $x_m, m \in S$ and let us introduce (K^{GM} and K^{GF} being positive constant parameters):

$$N_{m,n} = \begin{cases} \max(0, K^{GM} - M_m^n) & \text{if } m \in G^A \times G^B, n \in G^B \times G^A \\ \max(0, K^{GM} - M_n^m) & \text{if } m \in G^B \times G^A, n \in G^A \times G^B \\ K^{GF} & \text{if } m = n, m \in G^A \times \{\emptyset\} \\ K^{GF} & \text{if } m = n, m \in G^B \times \{\emptyset\} \\ 0 & \text{otherwise} \end{cases}$$

By construction, $N_{m,n}$ is symmetric and positive. $N_{m,n}$ is an affinity matrix because its values decrease when the costs/distances M_m^n (resp. M_n^m) increase. K^{GM} discards the shortcuts that are too expensive. The remaining ones are counted exactly two times. The parameter K^{GF} is used to discourage the matching.

In [139], according to the reformulation presented in [330], we define $y_m = \sqrt{x_m}$ - with component-wise square root - and consider the new matching problem:

$$\begin{aligned} & \max_y \sum_{m,n \in S} [N_{m,n}] y_m^2 y_n^2 \\ \text{s.t. } & \begin{cases} \sum_{m \in S_i} y_m^2 = 1 & \forall i \in G^A \\ \sum_{m \in S_j} y_m^2 = 1 & \forall j \in G^B \end{cases} \end{aligned}$$

This problem is a relaxation and a simplification of our initial problem (since uninteresting paths have been discarded) where the possibility of the absence of matching has been introduced and where non-matched segments are penalized by a cost depending on K^{GF} and K^{GM} . It is a forth-order tensor optimization that can be addressed directly [330]. Because the matrix N is sparse, spectral methods such as the power iteration method exploited in [330] are particularly interesting. This method - also called von Mises iteration method [342] - is used to determine the spectral radius of a matrix. When the considered matrix contains only positive components, according to the Perron-Frobenius theorem, this eigenvalue is positive, simple (the dimension of the its associated subspace is equal to 1) and its associated subspace is generated by a positive unit vector containing strictly positive components. In [343], this method has been extended into a high-order power iteration method working with tensors. The convergence of this procedure has been established [344], that provides also an another initialization method. This tensor power iteration method has recently been applied to solve graph matching problems [330]. In this paper, the tensor power iteration is presented as a simple way to deal with high-order matching potentials, such as the third order cliques that naturally appear when performing pose-invariant matching. There exists no guarantees on the convergence rate, because this rate is equal to the ratio between the two highest eigenvalues, that can be arbitrarily near to 1. In practice, however, the convergence is fast.

We determined the matchings ϕ and ψ by considering the boolean output that is the nearest to the real one provided by the power iteration output. Let us define Φ and Ψ by:

$$\begin{aligned} \Phi_{i,j} &= \phi_{i,j} V_\phi(i, j) & \forall (i, j) \in G^A \times G^B \\ \Psi_{k,l} &= \psi_{k,l} V_\psi(k, l) & \forall (k, l) \in G^B \times G^A \end{aligned}$$

We defined the new linking costs $\mathcal{C}(i, l)$ by:

$$\begin{aligned}\mathcal{C}_{i,l}^* &= \min \left(A_{i,l}, \min_{j,k \in G^B} [\Phi_{i,j} + C_{j,k} + \Psi_{k,l}] \right) \quad \forall (i, l) \in G^A \times G^A \\ \mathcal{C}(i, l) &= \min (\mathcal{C}_{i,l}^*, \mathcal{C}_{l,i}^*) \quad \forall (i, l) \in G^A \times G^A\end{aligned}$$

The first operation determines the best shortcut between i and l in G^A through G^B and compares it with a direct linking between i and l in G^A . The solution with lowest cost is retained, and symmetrized during the second step. Exploiting these new costs is of particular interest for linking the parts of the guide-wire that are partially invisible, but correctly matched to a strongly detected blood vessel.

4.4 Geometric-Iconic Graph-based Tracking

We adopt MRF-snake-spline approach describes the wire motion in a very compact manner. This method performs an iconic registration, while providing good solutions thanks to the MRF framework - under the assumption that the ideal displacements are part of the sampling exploited pattern.

The main limitation of this framework is that the size of the sampling pattern has to be increased to tackle large motions. As a result, the processing time is prohibiting. Moreover, large parts of the guide-wire can disappear, because of the noise and of motion blur. In these cases, tracking landmarks extracted along the wire is - temporarily - more reliable.

Aiming to take advantage of both approaches, we proposed an hybrid framework combining the iconic tracking [21, 167] and a geometric/landmark based tracking [345]. This latter is expressed as a matching problem that is expressed as an optimization problem coupling two solution spaces with supplementary MRF potentials. This coupling term is exact (see section 4.4.3) contrary to, for instance, [346, 347] that minimize an upper bound of our coupling terms. As a result, we have the guarantee that the optimization will not overestimate this term.

This section is devoted to the description of our method. After describing the improvements that we introduced into the iconic tracking [21] - in particular, by taking into account a wire orientation estimation - we describe the landmark extraction and the matching. The coupling potentials and a simple motion model based on templates are then presented and a coarse-to-fine procedure is finally explained.

Notations

The following notations hold in the whole section:

- $I(t)$ denotes the frame number t of the fluoroscopic sequence, and $I(\mathbf{x}, t)$ the intensity value at location $\mathbf{x}$

- $\mathcal{F}(\mathbf{x}, t) = (g(\mathbf{x}, t)\cos(\theta(\mathbf{x}, t)), g(\mathbf{x}, t)\sin(\theta(\mathbf{x}, t))) = (\mathcal{F}(\mathbf{x}, t)_x, \mathcal{F}(\mathbf{x}, t)_y)$ denotes the result of the steerable filter after the application of the Tensor Voting, and constitute the *wire support*. It combines a wire detector response $g(\mathbf{x}, t)$ and an estimation of the local wire orientation $\theta(\mathbf{x}, t)$.
- $c(i, t)$ for $i = 1 \dots N$ the location of the N control points describing the guide-wire conformation in the frame t .
- $\mathcal{P}$ denotes the discrete set of displacements of the control points.
- l_i is the label corresponding to the displacement of $c(i, t)$ and let $v(l_i)$ be the corresponding 2D displacement.
- without loss of generality, $s(c(i, t))$ is the curvilinear abscissa of the control point $c(i, t)$ - for uniform B-splines, these values range homogeneously from 0 to 1.
- $p(j, t)$ for $j = 1 \dots M$ the location of the M landmarks of the guide-wire at t . s_j denotes their curvilinear abscissa along the wire.
- $p^L(j, m, t+1)$ for $m = 1 \dots L$ the candidate number m extracted in the frame $t+1$, potentially matched with the landmark $p(j, t)$, L being thus the number of candidates extracted per landmark.
- $l_j = 1 \dots L$ the displacements $w(l_j)$ resulting from the matching: $w(l_j) = p^L(j, l_j, t+1) - p(j, t)$
- $\mathcal{X}(t) = (\{c(i, t)\}, \{p(j, t)\})$ formally denotes the configuration of the hybrid model of the guide-wire at time t .

With these notations, we can sum up our approach in saying that it estimates, for each frame I_{t+1} , the most probable $\mathcal{X}(t+1)$ knowing $\mathcal{X}(t)$ and the support $\mathcal{F}(t+1)$, by assuming that the $c(i, t+1)$ are part of the points obtained by applying the sampling pattern at the $c(i, t)$, and that the new location of the landmarks are part of the candidates extracted. In order to avoid a degeneracy due to the possible (long-time) displacements of the landmarks along the guide-wire, re-distribution/re-sampling of the landmark positions is frequently performed. The whole tracking is presented in the figure [Fig. (4.2)].

The Markov Random Field that is finally optimized contains one node per control point i , taking the label l_i , and one node per landmark j , taking the label l_j . It contains only singleton and pairwise potentials and its energy will be denoted by:

$$E_{MRF} = \sum_{y \in \{data, landmarks, \}} \mu_y \left(\sum_{k \in \{i, j\}} V_k^y(l_k) + \sum_{k, r \in \{i, j\}} V_{k, r}^y(l_k, l_r) \right)$$




Figure 4.2: Tracking procedure. The landmarks in the frame $t - 1$, $p(j, t - 1)$, are used when extracting the candidate in the frame t . In the same manner, the location of the control points $c(i, t - 1)$ in the frame $t - 1$ is employed for computing the admissible $c(i, t)$ according to the sampling pattern $\mathcal{P}$

where $\{i, j\}$ denotes all the control points and all the landmarks indices. As in [21], we will describe the guide-wire in the frame t with a cubic B-spline $\mathcal{C}(s, t)$ of the form:

$$\mathcal{C}(s, t) = \sum_i c(i, t) N_i(s, t)$$

where the N_i are the basis functions associated with the control points that, without loss of generality, could change from frame to frame (for instance, if the control points of the spline are re sampled), and where s is the curvilinear abscissa along the B-spline. $\mathcal{C}'(s, t)$ denotes the derivate of the curve. $\mathcal{C}_i(s, l_i, l_{i+1}, t)$ denotes the curve obtained when $c(i, t)$ is moved by $v(l_i)$ and $c(i + 1, t)$ by $v(l_{i+1})$, and $\mathcal{C}'_i(s, l_i, l_{i+1}, t)$ its derivative. We introduce the following *influence functions*:

$$N_{i,i+1}(s, t) = \frac{N_i(s, t) N_{i+1}(s, t)}{\sum_k N_k(s, t) N_{k+1}(s, t)} \quad \forall i < N - 1$$

that simply projects the influence of a given guide-wire point to the control points used to express/determine its displacement.

4.4.1 Iconic Tracking

The displacement of any point of a cubic B-spline is estimated using the displacements of the four neighboring control points. Therefore an exact iconic term would involve forth-order potentials, that are intractable. Following [21] we approximate the forth-order potentials with second order potentials by considering only the curve deformations that result from the displacement of the two control points influencing the most the location of each pixel: for each pair of control point, the potentials are computed according to the properties of the curve resulting

from their displacement, and weighted by the influence function. As a result, the potentials involved form a chain:

$$\sum_{k,r \in \{i,j\}} V_{k,r}^{data}(l_k, l_r) = \sum_{i=1}^{N-1} V_{i,i+1}^{data}(l_i, l_{i+1})$$

In order to obtain a robust data term, we took into account the local orientation of the spline. Without loss of generality, the pairwise potentials are thus given by:

$$V_{i,i+1}^{data}(l_i, l_{i+1}) = \int_{s(c(i,t))}^{s(c(i+1,t))} N_{i,i+1}(s, t) \Phi(\mathcal{F}(s, t), \mathcal{C}_i(s, l_i, l_{i+1}, t), \mathcal{C}'_i(s, l_i, l_{i+1}, t)) ds$$

where $\Phi(., ., .)$ decreases when the curve passes through pixels that are likely to be part of the wire (where $g(\mathbf{x}, t)$ is high) and where the spline derivative is coherent with the visual support orientation, $\theta(\mathbf{x}, t)$.

More precisely, we applied the idea presented in [26] by choosing a cost that diminishes when the scalar product $\langle ., . \rangle$ between the curve derivative and the oriented support of the wire increases. We shall emphasize that the scalar product has to be normalized by the norm of the derivative of the curve to be independent from the wire parametrization. We finally used the following potentials:

$$\begin{aligned} V_{i,i+1}^{data}(l_i, l_{i+1}) &= \int_{s(c(i,t))}^{s(c(i+1,t))} N_{i,i+1}(s, t) \psi \left(\frac{|\langle \mathcal{C}'_i(s, l_i, l_{i+1}, t), \mathcal{F}(\mathcal{C}_i(s, l_i, l_{i+1}, t), t) \rangle|}{\|\mathcal{C}'_i(s, l_i, l_{i+1}, t)\|} \right) ds \\ &= \int_{s(c(i,t))}^{s(c(i+1,t))} N_{i,i+1}(s, t) \psi(g(\mathcal{C}_i(s, l_i, l_{i+1}, t), t) |\cos(\phi(\mathcal{C}_i(s, l_i, l_{i+1}, t)))|) ds \end{aligned}$$

where $\phi(\mathcal{C}_i(s, l_i, l_{i+1}, t))$ is the angle between the derivative $\mathcal{C}'_i(s, l_i, l_{i+1}, t)$ and the unitary vector of orientation $\theta(\mathcal{C}_i(s, l_i, l_{i+1}, t), t)$, and where (γ being a positive constant):

$$\psi(x) = e^{-\gamma x}$$

These potentials increase when the response of the detector along the spline diminishes and when the spline tends to become orthogonal to the orientation determined by the detector. The only way to reduce them consists in determining a spline that passes through locations where the detection is strong and is simultaneously coherent with the detected local orientation.

4.4.2 Landmark Matching

In this section, geometric tracking towards improving tracking performances is described. We first describe the landmark extraction process. Then, the choice of the candidates for the matching of each landmark is presented. The graph-based matching problem solved into our MRF framework is finally presented.

Our method extracts reliable landmarks by sampling points along the wire and retaining only the locations where strong and coherent wire detections accumulate. These locations are determined by computing the main eigenvalue of the following tensors:

$$T(\mathbf{x}, t) = \sum_{\mathbf{y}} e^{-\frac{\|\mathbf{x}-\mathbf{y}\|^2}{2\sigma_L^2}} \begin{pmatrix} \mathcal{F}(\mathbf{y}, t)_x^2 & \mathcal{F}(\mathbf{y}, t)_x \mathcal{F}(\mathbf{y}, t)_y \\ \mathcal{F}(\mathbf{y}, t)_x \mathcal{F}(\mathbf{y}, t)_y & \mathcal{F}(\mathbf{y}, t)_y^2 \end{pmatrix}$$

These tensors are inspired from the structure tensor [38], but they rely on the wire support $\mathcal{F}(t)$. This support is more reliable than the image intensity and the steerable filter response, because it has been regularized by Tensor Voting.

For each landmark $p(j, t-1)$ extracted in the frame $t-1$, we seek L candidates location $p^L(j, m, t)$ in the current frame $I(t)$ that are similar to $p(j, t-1)$ according to a given metric - that is related to a group of allowed transformations. Aiming, like in [345], to penalize all the deformations of the wire except the translations, we use the following affinity score to decide if a location $\mathbf{y}$ is a good candidate to match with $p(j, t-1)$:

$$\mathcal{M}_j(\mathbf{y}) = \langle T(\mathbf{y}, t) | T(p(j, t), t) \rangle$$

where $\langle . | . \rangle$ denotes the inner product between matrices. We must emphasize that this choice was dictated by efficiency consideration (tensors computations are performed only once per frame) and does not correspond to a very discriminative matching criterion. A better criterion could exploit the normalized cross-correlation of patches centered around $p(j, t)$ and $\mathbf{y}$. Such a choice would allow contrast intensity changes. Another choice could rely on the extraction and the comparison of features and descriptors from these patches such as affine invariant regions and SIFT [36, 37, 314, 315, 316].

In order to deal with missing matches, we introduce an additional label indicating that no candidate is sufficiently interesting to be matched to $p(j, t-1)$. We penalized the choice of this label by introducing the following singleton potentials [345]:

$$V_j^{landmarks}(l_j) = \begin{cases} K & \text{if } l_j = L+1 \\ 0 & \text{otherwise} \end{cases}$$

The penalization of the deformation of the landmarks configuration was introduced by the following pairwise potentials, that were considered for each pair of landmarks (j, k) (Γ being a constant positive threshold):

$$\begin{aligned} V_{j,k}^{landmarks}(l_j, l_k) &= \begin{cases} 0 & \text{if } l_j = L+1 \text{ or } l_k = L+1 \\ \min(\|\mathbf{u}(j, l_j, k, l_k) - \mathbf{u}(j, k)\|^2, \Gamma) & \text{otherwise} \end{cases} \\ \mathbf{u}(j, l_j, k, l_k) &= p^L(j, l_j, t) - p^L(k, l_k, t) \\ \mathbf{u}(j, k) &= p(j, t-1) - p(k, t-1) \end{aligned}$$

These potentials preserve the distance between landmarks and the orientation of the pairs of landmarks. They are obtained by applying the robust estimator:

$$x \mapsto \min(x^2, \Gamma)$$

to the penalty [348], that was exploited in the context of brain MRI registration. As a result, our potentials are more robust than the ones in [348]. In particular, the threshold avoids a predominance of the potentials related to distant landmarks. In accordance with the candidate extraction, this choice does not penalize the wire translations.

4.4.3 Coupling Part

In order to couple the iconic and the geometric tracking, we introduce potentials that penalize curve deformations producing correspondences that do not correspond to the landmark matching [346, 349, 347]. Let us consider that Gaussian law expresses the error of landmark localization due to the presence of noise, then minimizing the minus log-likelihood of the displacements and the matching knowing the wire support leads to minimize, for each landmark j , a tracking consistency cost proportional to:

$$\epsilon_j = \left\| \sum_i N_i(s_j, t) v_i(l_i) - w_j(l_j) \right\|_2^2$$

because the displacement $\sum_i N_i(s_j, t) \mathbf{v}_i(l_i)$ is the landmark displacement described by the displacement of the control points, and $w_j(l_j)$ the displacement resulting from the matching. These costs involve fifth-order cliques (containing one landmark and the four control point describing its location) but this expression boils down to:

$$\begin{aligned} \epsilon_j &= \left\| \sum_i N_i(s_j, t) v_i(l_i) - w_j(l_j) \right\|_2^2 \\ &= \sum_{i \neq k} N_i(s_j, t) N_k(s_j, t) \langle v_i(l_i), v_k(l_k) \rangle - 2 \sum_i N_i(s_j, t) \langle v_i(l_i), w_j(l_j) \rangle \\ &\quad + \sum_i N_i(s_j, t)^2 \|v_i(l_i)\|_2^2 + \|w_j(l_j)\|_2^2 \end{aligned}$$

where $\langle, \rangle$ denotes the standard scalar product. This expression contains only pairwise potentials and singleton potentials, and can therefore be efficiently handled by Markov Random Field solvers.

In spite of its mathematical correctness, this model is incomplete in practice, because many different displacements of the control points can produce the observed landmark motion. It is thus necessary to introduce potentials acting as a curve rigidity. A simple way to do this consists in penalizing the square Euclidean norm of the difference between the motions of pairs of control points. We introduced thus the following pairwise potentials in order to enforce the rigidity of the tracker around the landmark (where $\mathcal{X}(j)$ contains the index of the control points describing the location of $p(j, t)$):

$$r_j = \sum_{i \in \mathcal{X}(j)} \|v(l_i) - v(l_{i+1})\|_2^2$$

As a result, the potential coupling the iconic and the geometric tracking are the following, where the parameter δ governs the rigidity of the tracker:

$$\sum_k V_k^{coupling}(l_k) + \sum_{k,r} V_{k,r}^{coupling}(l_k, l_r) = \sum_j \epsilon_j + \delta \sum_j r_j$$

The hybrid tracking presented so far addresses most of the challenges that we have described. However, when the blur caused by large motion makes the thinnest parts of the wire invisible during a few frames, neither the geometric nor the iconic tracking can prevent a deformation of the spline parts that are tracking them. This is due to the fact that landmarks are not reliable for the tracking of thin wire parts, and that the wire support exploited by the iconic tracking is temporarily useless. In order to address this challenge, we have enforced the temporal consistency of the tracking as explained in the next section.

4.4.4 Dynamic Consistency

Large wire displacements often create a motion blur sufficient to make long parts of the wires invisible and data terms become unreliable. In order to make the tracking more robust with regard to this problem, we introduced a dynamical model-like behavior in the process through potentials locally penalizing the change of the shape, like snake-length preserving internal energy does [29, 21, 167]. Similarly to the data term, these potentials are expressed as an integral along the curve, that is decomposed into pairwise terms forming a chain:

$$\sum_{k,r \in \{i,j\}} V_{k,r}^{template}(l_k, l_r) = \sum_{i=1}^{N-1} V_{i,i+1}^{template}(l_i, l_{i+1})$$

Let $\mathcal{T}$ be a template of the curve C that is updated following an exponential forgetting law:

$$\mathcal{T}(s) \leftarrow \frac{memory - 1}{memory} \mathcal{T}(s) + \frac{1}{memory} C(s)$$

where the parameter *memory* is related to the rate of forgetting. We adopted a variant of the local length preserving potentials [21] for exploiting this model. With our notations, the length-preserving potentials [21] are given by:

$$V_{i,i+1}^{template}(l_i, l_{i+1})^{[21]} = \int_{s(c(i,t))}^{s(ci+1,t))} \frac{(\|C'(s, l_i, l_{i+1}, t)\| - \|\mathcal{T}'(s)\|)^2}{\|\mathcal{T}'(s)\|^2} ds$$

In order to penalize the curves rotation in addition to the local dilation/shrinkage, we considered the variant:

$$V_{i,i+1}^{template}(l_i, l_{i+1}) = \int_{s(c(i,t))}^{s(ci+1,t))} \frac{\|\eta C'(s, l_i, l_{i+1}, t) - \mathcal{T}'(s)\|^2}{\|\mathcal{T}'(s)\|^2} ds$$

where η - chosen slightly smaller than 1 - introduces a ballooning force. This model is more constraining than the length preserving terms exploited in [29, 21], because it penalizes the orientation changes of the derivative of the curve.

4.5 Experimental Validation

4.5.1 Segmentation through Graph/Image Fusion

We validated the exploitation of image fusion for guide-wire segmentation with our multiple-structures segmentation pipeline. In order to reduce the complexity of the linking, whose iterations requires $O(n^4)$ operations for n primitives, we introduced a preprocessing step before the linking [139]. This component aims at rejecting isolated primitives and group line-segments that should obviously be considered in the same manner.

For each intervention, one fluoroscopic image and one contrast-enhanced image corresponding to the same moment in the cardiac cycle were considered. The endpoints of the most interesting wire were provided to the algorithm. Results were measured in terms of missed detection rate (MDR) and false detections rate (FDR) at 5 pixels (1mm) [133, 24]. In order to take the presence of multiple structures into account, we have also annotated other structures of interest (other wires, catheters) and compared all the structures segmented with this ground truth, obtaining MDR^{all} , FDR^{all} .

The graph G^A was defined by:

$$A_{kl} = \mu(\beta_1 d^s(k, l) + \beta_2 e^s(k, l)) \quad \forall k, l \in G^A$$

where d^s is the Euclidean distance between the second extremity of k and the first extremity of l , where e^s is the scale invariant Elastica [271] corresponding to this link and $\mu(\cdot)$ discards edges with costs higher than a threshold K^{GM} :

$$\mu(x) = \begin{cases} x & \text{if } x \leq K^{GM} \\ \infty & \text{otherwise} \end{cases}$$

G^B was defined in the same manner and we chose transition costs V_ϕ and V_ψ :

$$V_\phi(i, j) = V_\psi(j, i) = \gamma_1 d^m(i, j) + \gamma_2 (1 - \cos(i, j))$$

where d^m denotes the Euclidean distance between the centers of the segments i and j , $\cos(i, j)$ the cosinus of the angle between them and $(\gamma_1, \gamma_2) = (0.5, 100)$. We used the following segment cost, where $\mathcal{L}(i)$ denotes the length of i : $\mathcal{C}(i) = -0.5 \mathcal{L}(i)$. We extracted 1400 points for an image of size 512×512 (adapting this number to the area of the image) from a grid of step 3 for fluoroscopic images and of step 5 for enhanced images. The guide-wires were extracted at scale 2 and the blood vessels at scales $\{3, 5\}$. For line-segment extraction, we set $(\alpha, K^{MCE}, d_\emptyset) = (10, 275, 150)$. $(\beta_1, \beta_2, K^{GF}) = (0.01, 30, 185)$ defined the travel costs through the graphs G^A and

	MDR	FDR	MDR^{all}	FDR^{all}
points	0.128	0.899	0.17	0.858
line-segments	0.246	0.862	0.299	0.818
line-segments chosen by the preliminary linking	0.274	0.667	0.334	0.578
final wire(s)	0.275	0.243	0.323	0.36
final wire(s) without preliminary linking	0.467	0.415	0.448	0.501

Table 4.1: Mean results of our approach. **MDR** is the proportion of ground truth pixels lying at more than 5 pixels of the segmented wires and **FDR** the proportion of segmented structures lying at more than 5 pixels of the ground truth. All the finally segmented structures are used for measuring MDR^{all} and FDR^{all} , whereas only the first (constrained) wire is used for measuring **MDR** and **FDR**. The preliminary linking reduced the cumulated computational time over all the 48 test sequences from 15100 to 19 seconds.




Figure 4.3: The four different evaluations of our complete method (blue), without image fusion (green), with image fusion but single wire delineation [133] (red), without image fusion and with [133] (black).




Figure 4.4: **On the first row:** the fluoroscopic image, the line segment extracted from it, the matching Φ (line-segments from the fluoroscopic image, in green, are matched with blood vessel centerlines parts, in blue) and the output of the preliminary linking. **On the second row:** the contrast-enhanced image, the line-segment extracted from it, the matching Ψ (between green centerlines and blue fluoroscopic line-segments) and the structures finally segmented

G^B . $K^{GM} = 200$ was used to determine the neighbors of a segment in these graphs, and K^{MSS} was set to 25.

The table [Tab. (4.3)] summarizes the quantitative results obtained on 48 images (extracted from 48 sequences) of the database $D2$ (made of 80 8-bits fluoroscopic and contrast-enhanced sequences). The first figure [Fig. (4.4)] illustrates the main segmentation step, while [Fig. (4.5)] emphasizes the link modification resulting from the incorporation of the information originated from the contrast-enhanced images. All these results confirm that the image fusion improves the results. Quantitative comparison are provided in the figure [Fig. (4.3)], demonstrating that preliminary linking increases the performances (in addition to reducing the complexity, as indicated in the table Tab. (4.3)).

4.5.2 Tracking

Because the number of pairwise term estimations, that constitutes the bottleneck of the method, increases with the number of labels, reducing the size of the set of admissible displacements seems to be crucial for speeding up the algorithm. For this reason, we adopted a coarse-to-fine approach [345]: we performed T optimization step (instead of only one), indexed with $\tau = 1 \dots T$, with a sampling pattern $\mathcal{P}^\tau$ having less layers and a scale σ^τ decreasing with τ (where a is a constant parameter):

$$\sigma^\tau = 2^{(a-\tau)}$$




Figure 4.5: Qualitative effects of the exploitation of the contrast-enhanced image. Depending on the rows: (i) line-segments extracted in the fluoroscopic image (ii) costs reduced by the graph fusion (in red) (iii) structures finally segmented

Such an approach is similar to the one presented in [339] and very popular in computer vision.

The database used to validate our tracking method is made of 20 long fluoroscopic sequences containing each more than 200 images. This length is long with respect to the sequences currently handled in clinical settings and with respect to other recent publications on the domain [24, 21]. The frame rate of this sequences is equal to 10 images per second. These sequences were chosen to span as many clinical situations as possible. The size of the frames varied from 512×512 to 1000×1000 pixels.

Three sequences were used to calibrate the model parameters [345]. The tensor voting was applied at scale 4.5 on the response of the steerable filter [64] of order 2, $\mu = 0.0$ and scale 1.8. We set $\mu_{data} = \mu_{template} = 1.0$, $\mu_{landmarks} = 2.10^{-4}$, $\mu_{coupling} = 3.10^{-5}$, $\delta = 0.1$, $\epsilon = 0.9$, $K = 45000$ for the landmarks not matched and displacements steps were equal to $(25.0)2^{-\tau}$, $\tau \in \{0, \dots, 4\}$. We detected landmarks with $\sigma_L = 2.5$. The displacements pattern described in [21] was used to sample 25 displacements (see the figure [Fig. (4.6)]). We set $L = 24$, $\gamma = 6.10^{-3}$, $\Gamma = 500$ and *memory* = 7.

The remaining 17 sequences were used to perform the validation. Two metrics were used: in addition to the Missed Detection Rate/False Detection Rate metric, the Euclidean distance between the tip of the guide-wire and the tip of the tracker were also measured.

During the validation, the performances were measured every twenty-five frame. The figure [Fig. (4.6)] presents the quantitative results obtained. The figure [Fig. (4.7)] presents the evolution of the missed detection rate during the tracking, and compares the results obtained with our method with the results obtained with [21].

Our method outperformed [21], due to two reasons: the visual support that we extract is more complete, and our tracking data term takes the local orientation estimation into account and because of the use of a template and the landmarks matching, that make our tracker able to follow large deformations.

4.5.3 Segmentation and Tracking

Given that our original purpose was to provide a fully automatic tracking method, we tested if the output of our segmentation method was sufficient for being used as an input of our tracking methods. In order to do so, we applied a semi-automatic variant of our segmentation method to the first frame of the fluoroscopic sequences of our tracking database, and we used this initialization instead of the ground truth used until now. This variant requests the user to indicate a location near to the beginning of the wire and a location near to its end. The same wire support is extracted for the segmentation and the tracking: the steerable filters responses regularized by tensor voting.

For the segmentation, we tiled our images and considered only a small set of points from each tile. These points were the 120 points with strongest responses taken from a grid of step 4, and the size of the tiles was 200×200 points during our experiments. We used the same guide-wire detector as in the previous section. A




Figure 4.6: (i) the sparse sampling pattern and the labels attributed to it (ii) mean tracking results (over the 17 validation sequences) estimated every 25 frame. The sequence of result (a) was obtained by applying the complete method [21] on our database. The result (b) is the result reported in [24]. It is only presented as an illustration (because it was obtained with a different database)




Figure 4.7: Missed Detection Rates measured every 25 frames for the 17 validation fluoroscopic sequences, with the tracking [21] (blue boxes) or with our method (black boxes)

distance exploiting 100.0 times the Elastica, a cluster penalty equal to 15000.0 and a background pixels cost of 16000.0 were used during the clustering. The MSAC was performed at scale 6.0. For the linking part, the following linking costs and line-segment rejection costs were used:

$$\begin{aligned} C(\mathbf{x}, \mathbf{y}) &= \delta_{iconic} \int_{\mathbf{x}}^{\mathbf{y}} \psi \left(\frac{|\langle \mathcal{F}(\mathbf{s}, 0), \mathbf{y} - \mathbf{x} \rangle|}{\|\mathbf{y} - \mathbf{x}\|} \right) ds \\ &\quad + \delta_{euclidean} \|\mathbf{y} - \mathbf{x}\| + \delta_{elastica} [e(\mathbf{x}, \mathbf{y})] \\ \mathcal{C}(s_n) &= \int_{\mathbf{a}_n}^{\mathbf{b}_n} 1 - \psi \left(\frac{|\langle \mathcal{F}(\mathbf{s}, 0), \mathbf{b}_n - \mathbf{a}_n \rangle|}{\|\mathbf{b}_n - \mathbf{a}_n\|} \right) ds \end{aligned}$$

with $\delta_{iconic} = 6.5$, $\delta_{elastica} = 1.8$ and a small $\delta_{euclidean} = 0.1$ and $\psi(x) = e^{-\beta x}$, $\beta = 6.10^{-3}$ and any line-segment $s_n = (\mathbf{a}_n, \mathbf{b}_n)$. The "iconic" part of the linking cost ensures that the link is compatible with the guide-wire support and takes the local orientation estimation into account. We have assumed with this choice of costs that all the optimum links are straight lines. For the tracking part, the parameters were the same as in the previous section.

The figure [Fig. (4.8)] presents the extracted support, the line-segment obtained and the final guide-wire determined for five sequences. Quantitative results are provided by the figures [Fig. (4.9,4.10)], that demonstrate that the performances decrease that is measured when the semi-automatic segmentation method is used as an input for the tracking procedure, is noticeable but limited.

4.5.4 Limitations of the Tip-distance Metric

Since the tip of the guide-wire is made of a more opaque material, this measure is not well designed: long distances can correspond to a wire that has not been extended to fit the whole extracted support, and this situation is not critical because the tip of the wire is easy to see/detect. On the contrary, a small distance orthogonal to the true wire position could actually describe a severe failure of the tracking. It would be more interesting to measure a distance orthogonal to the true wire location and report this distance with a missed detection rate during future work.

4.6 Conclusions and Future Work

In this chapter, we have first proposed a graph matching based method that fuses a set of line-segments delineating the blood vessel centerlines detected in a contrast-enhanced image with the ones extracted in fluoroscopic images. Such a matching is exploited for improving the linking costs that are used for delineating the guide-wires. We have then presented an hybrid method dedicated to the tracking of highly non-rigid and moving curvilinear structures similar to guide-wires. This method involves an iconic part that is linked to a geometric part by a coupling. These models are considered simultaneously in a single-shot Markov Random Field




Figure 4.8: (a) input images (b) regularized responses (intensity has been windowed for increasing the contrast) (c) segmented line-segments (d) segmented guide-wires




Figure 4.9: (a) comparison of mean tracking result obtained with the ground truth initialization and with the semi-automatic initialization. (i) was obtained with [21], and (ii) is a recall of [24] (these results are already presented in [Fig. (4.6)]). (b) tip tracking error in pixel, for method [21] (blue) and for our method when initialized with ground truth (black) or segmentation output (green)




Figure 4.10: (a) comparison of Missed Detection Rates obtained during the tracking with ground truth initialization (in black) and with semi-automatic initialization (in blue). (b) related False Detection Rates obtained with ground truth at left and with semi-automatic initialization (at right).

framework, that introduces interaction between them and provides robust results. The iconic representation is based on a cubic B-spline and is made robust (with respect to noise masking the wire support) by incorporating a local similarity criterion with respect to a template of the structure, that is updated during the tracking. The geometric part consists in the detection of reliable landmarks along the wire, that are matched with their potential new locations in a way that globally preserves the shape of the structure. Contrary to previous work [346, 349, 347] on hybrid geometric-iconic registration, the coupling term proposed here is not an upper bound of the distance between the displacement/distortion fields described by the iconic and the geometric part. It corresponds on the contrary to the exact value of the square Euclidean distance between them.

Let us briefly describe three improvements of the existing modules that could be easily investigated. Firstly, the proposed landmark extraction method extracts points belonging to a strongly detected part of the wire. However, the criterion exploited in this work do not insure that the points finally extracted are reliably different from their neighbors. Introducing a saliency measure at this step could therefore increase the precision of the method. Exploiting a corner detector is an alternative way to explore.

Secondly, the presented method penalizes rotations of the wires, both at the level of the candidates extraction - that tends to extract candidate with similar orientations - and at the MRF level. The set of rotations could easily be introduced into the admitted transformations by exploiting at the first level one of the numerous approaches performing orientation-independent comparison of patches that exist, such as comparing SIFT points [36, 37] ; and at the second level by penalizing in an other manner the inter-landmarks distance modifications. As a result, the tracker would globally rotate at no cost. It is however not clear if these modifications will improve the results by reducing the penalty considered for the real motions, or if they will deteriorate them by increasing too much the set of weakly penalized deformations.

Finally, a better sampling strategy could be exploited. The coarse-to-fine strategy aiming to reduce the computational burden did not performed as well as expected: the computational time saved at each optimization step by the reduction of label set was not sufficient to compensate the burden introduced by the calculus of new singleton potentials and by the supplementary iterations. A better way to reduce this set could consist in applying a generalized Hough transform [350] to the set of candidate landmarks. Because they are supposed to lie on the target location of the guide-wire, considering only the displacements that are compatible with the most landmarks would allow to dramatically reduce the displacement considered without performing a coarse-to-fine process. A MSAC [18] procedure could also be envisaged to select the admissible displacements: pairs of control points displacements would be picked randomly. The score of these pairs would depend on the number of candidates lying near to the spline obtained after this displacement. In this manner, the displacements offering the most possibilities for the landmark matching would be selected.

Among the various extensions of our work, the four following ones constitute interesting future work.

Improving the existing modules, for instance by exploiting a better rigidity criterion, a better landmarks detector and better optimization parameters (sampling pattern, coarse-to-fine strategy ...) - that influence significantly the complexity - appear as the first extension to try.

In a similar manner, other graph matching methods could be considered. Graph matching being the subject of a rich literature, interesting approaches could probably be applied to our problem. The problem of partial matching would remain: because primitives are entirely matched in our framework, the result can be influenced by relatively small deformations of the blood vessel modifying the number of extracted primitives. This last problem could be addressed by splitting the primitives into smaller pieces of constant length. It has been addressed in [351], that performs *fractional* assignments between sets of line-segments. [352], that matches shock trees extracted from objects shapes could also be exploited, since the coronary arteries also form a tree that can easily be extracted from the shape of the blood vessels wall (these latter being easily extracted with a simple gradient operator). We refer the reader to the review [353], that describes many interesting graph matching methods.

Incorporating a guide-wire displacement model, such as the motion simulation in [354] could probably constrain the deformations to be more realistic.

Studying a dense motion model would be interesting. The wire displacements are indeed difficult to establish because this motion combines cardiac and respiratory motion with the displacements controlled by the cardiologist. A *dense* motion model would on the contrary only capture the biological motion and would therefore be easier to describe. The limitations of this approach are threefold: first, the combination of cardiac and respiratory motion is still difficult to predict ; second, it is difficult to reconstruct a motion model from the fluoroscopic images, because they result from the 2D projection of a 3D multi-layer transparent structure ; finally, one should develop a coupling between this model and the guide-wire displacement allowing the wire to freely move into the artery where it is pushed.

Chapter 5

Review of Contributions and Future Work

5.1 Review of contributions

This work addresses the detection, the segmentation and the tracking of the guide-wires in fluoroscopic imaging. The detection is particularly challenging because of the low signal-to-noise ratio of the fluoroscopic sequences and because the wires are thin. Many confusing structures such as staples and rib borders are also visible. The tracking is particularly challenging due to the motion blur, the injection of contrast agent and the manipulation of the wire during the intervention, and because the cardiac and respiratory motion cause large wire deformations.

We introduced two guide-wires detection approaches, one that is based on modern machine learning methods and the second that is based on adaptative filter selection. For guide-wires segmentation, we propose a bottom-up pipeline that extracts a set of points according to the guide wire detector, that groups these points into local segments (clusters of points), extracts the characteristic geometric primitive per cluster and that links these primitives in order to delineate the wires. Two mathematical formulations/frameworks are considered for linking, one that is exploring linear programming/MRFs and the second that is based on a meta-heuristic step. The first approach consists of linking the line segments towards complete reconstruction of the guide-wire and in this context methods of increasing degree of interactions between segments are considered and solved through global optimization method. Due to the limited number of segments, we consider as well a meta-heuristic local search formulation that also performs ordering as an alternative to the global optimization methods. Since the guide-wires lie into the blood vessels, we also study an image fusion incorporating the information provided by the segmentation of the blood vessels that are visible in a contrast-injected image of the same patient. This is achieved by formulating a combined segmentation/graph-matching registration framework that performs segmentation and matching between the two complementary modalities. To this end, we explore forward and backward transitions between the line segments of the two modalities towards establishing their correspondences. This complex optimization problem is solved using a recent tensor based method. Last but not least, we propose a new robust guide-wire tracking combining the iconic tracking method [21, 32, 33] with a landmark based tracking. This method also explores a combined segmentation/registration/graph-matching approach where the aim is to find guide-wire optimal paths in a new image while establishing correspondences with salient points of the guide-wire between two successive frames. The method is endowed with a simple dynamical model that also imposes certain consistency between successive geometric forms of the guide-wire.

Promising experimental results throughout the entire chain of processing, as well as for the overall pipeline demonstrate the potentials of our contributions.

5.2 Review of Future Work

Interesting future directions to investigate were presented at the end of each chapter. Instead of citing them again, we would like to stress the transversal ideas that can be summarized. We first detail extensions of the method towards clinical use and then focus on potential theoretical extensions that could have an impact going beyond the clinical context considered in this thesis.

Firstly, all the components that were proposed in this thesis could be investigated again one by one. The modularity of the pipelines that were presented for the segmentation and the tracking allow to test multiple variants in a transparent way. Secondly, it would be interesting to study the possible integration of the different steps of the methods that were presented here, because a reduced number of modules would increase robustness and a reduced number of parameters to set up. The integration has already been done for the tracking, since the iconic and the geometric tracking that are combined in our method are expressed (and coupled) in a unique Markov Random Field. It is however perhaps possible to reduce the number of potentials that are involved in this model, in order to reduce the number of parameters of the approach. Thirdly, the detection is fundamental for all the tasks, but it seems very difficult to obtain satisfying results in that domain. The segmentation requires indeed the best possible detection but there is a strong trade-off between time and performances. Moreover, we were not able to obtain very good detection results: we obtained at most a true positive rate of 35 – 40% (for an admissible false positive rate of a few percents) , what is poor compared to the 80 – 90% that are commonly obtained in the machine learning literature for the applications that are qualified as *successful* or *promising*. The detection approaches exploited for the tracking, that need to be dramatically faster and more robust, naturally exhibit even lower performances. Having said that, it seems us crucial to try to develop a novel guide-wire detector, but we are not sure that the guide-wire detection can be tackled in a satisfying fully automatic and autonomous manner. The segmentation and tracking methods applied afterwards will certainly have to be robust enough to discard the numerous outliers, distinguish the confusing structures from the interesting ones, and overcome the missing detections. Towards clinical setting, one could imagine to develop massively parallel guide-wires detectors implemented on GPU, while increasing the ability of the higher-level components of rejecting outliers. We also think that the discrimination between the guide-wires and the most confusing structures, such as the rib borders, will only be possible at the highest level (for instance, by segmenting the catheters) with the help of a human user or by exploiting clinical context informations, such as the angle at which the images are acquired.

From theoretical view point, the proposed formulations explore mostly what is so called pairwise constraints between line-segments when reconstructing the guide-wire. Higher order potentials could be of great interest within such a context for the following reasons: (i) increase of robustness due to the ability to impose consistency regarding the geometry of the guide-wire over arbitrary subsets of segments, (ii)

invariance with respect to global linear geometric transformations due to the fact that the relative linear motion can be compensated when considering interactions going beyond three points, (iii) exact modeling and exact satisfaction of desired geometric properties like curvature, B-spline interpolation, etc. Last but not least, the use of higher order graphs allows us to learn optimal methods of representing the observed guide-wire, or in simple words finding the minimal number of higher order cliques that could optimally encode the geometric information of the guide-wire.

Furthermore, one of the main limitations of the proposed formulation is the lack of modeling the dynamical behavior of the guide-wire. We have consider very simplistic priors in tracking based on rather short modeling of the behavior of the guide-wire. The observed motion and deformation are due to three factors, heart motion, respiratory motion and deformation of the guide-wire due to the surgeon. The first two components of the motion can be considered to certain extend periodic and therefore could be modeled using a dynamical system. The last component of the deformation is "limited" and therefore the entire deformation sequence can be approximated using a sparsity constraint, that in theory should be able to provide the decomposition of the observed deformation into the dominant/periodic and the local deformation components on-line. Such a model can then be used as a prior in tracking, and could have a substantial impact on the temporal performance of the method.

Appendix A

Demonstrations

A.1 Multiple Guide-wires Segmentation is NP-Hard

One can increase the costs $\mathcal{C}(i)$ for considering a given primitive i as an outlier until no primitive is rejected. If one choose $\mathcal{C}(i) = K^{MSS}$, it is indeed always more interesting to create a guide-wire containing i than rejecting i as an outlier.

For this reason, the partition of the set of primitives into a spanning set of chains is a singular case of the multiple guide-wires segmentation problem that we consider in this thesis. By showing that a sub-problem of this specific multiple guide-wires segmentation problem is NP-hard, we prove in the sequel that the generic segmentation of multiple guide-wires is NP-hard.

Lemma

When no primitive is rejected and when the linking costs are independent on the primitive orientation, the segmentation problem is NP-hard.

Proof

We prove this lemma by constructing an undirected graph and by showing that, in the specific case described, the multiple guide-wires segmentation is an instance of **Degree Constrained Minimal Spanning Tree** (DCMST) on this graph. Because the DCMST is known to be NP-hard [355, 356], this concludes the proof.

Since one assumed that no primitive has been rejected, the costs $\mathcal{C}(i)$ are useless and we discard them. Let assume that the linking costs do not depend on the primitive orientation (what happens, for instance, when the considered primitives are points). Then (denoting with $\mathcal{P}$ the set of oriented primitives and $\bar{i}$ the oriented primitive with same support and inverse orientation as i):

$$\begin{aligned}\mathcal{C}(i, j) &= \mathcal{C}(i, \bar{j}) = \mathcal{C}(\bar{i}, j) = \mathcal{C}(\bar{i}, \bar{j}) \\ &= \mathcal{C}(\bar{j}, \bar{i}) = \mathcal{C}(j, \bar{i}) = \mathcal{C}(\bar{j}, i) = \mathcal{C}(j, i) \quad \forall (i, j) \in \mathcal{P}\end{aligned}$$

Let then build an undirected graph G containing one node $\textcircled{i}$ per pair of oriented primitives $\{i, \bar{i}\}$, and one edge per existing edge in the graph built by our linking method, weighted by:

$$w(\textcircled{i}, \textcircled{j}) = \mathcal{C}(i, j)$$

Let also introduce a supplementary node ξ linked with every other node by edges of weights:

$$w(\xi, \textcircled{i}) = K^{MSS}$$

This construction is illustrated by the figure [Fig. (A.1)]. Finally, let define the degree constraints $(\mathcal{D})$:




Figure A.1: (i) Illustration of the graph construction (ii) the solution is similar to this tree of root ξ (where the link penalized by K^{MSS} are dashed), and is therefore equivalent to the solution of the original multiple guide-wires segmentation problem.

- the degree of ξ is not constrained (or equivalently: constrained to be smaller than the number of nodes of the graph)
- the degree of each other node is constrained to be smaller than 2.

Then, the multiple guide-wires segmentation is equivalent to the DCMST on G weighted by w under the degree constraints ($\mathcal{D}$).

The degree constraints ensure indeed that the solution $\mathcal{S}$ of the DCMST on G is made of chains that are linked together in ξ and that partition the original set of primitives. The cost for $\mathcal{S}$ is equal to the sum of the linking costs between the successive elements of the chains plus K^{MSS} times the number of chains created. This cost is equal to the cost considered by our segmentation method. Thus, the two problems have the same solution. ■

Bibliography

- [1] e. a. Robert Beaglehole, *The World Health Report 2004 Changing History*, 2004, p. 122.
- [2] the American Heart Association Statistics Committee and S. S. Subcommittee, “Heart disease and stroke statistics – 2012 update : A report from the american heart association,” *Circulation*, vol. 125, pp. e2–e220S, 2011.
- [3] C. Dotter and M. Judkins, “Transluminal treatment of arteriosclerotic obstruction.” *Circulation*, vol. 30, pp. 654–670, November 1964.
- [4] F. Williams, “A method for morefully determining the outline of the heart by means of the fluoroscope, together with other uses for this instrument in medicine.” *Boston Medical Surgery Journal*, vol. 135, pp. 335–337, 1896.
- [5] 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*, vol. 4, pp. 565–576, 2011.
- [6] V. Bismuth, “Image processing algorithms for the visualization of interventional devices in x-ray fluoroscopy,” Ph.D. dissertation, ESIEE Paris, 2012.
- [7] 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*, vol. 13, pp. 819–845, 2009.
- [8] M. S. Mabrouk, N. H. Solouma, and Y. M. Kadah, “Survey of retinal image segmentation and registration,” *GVIP Journal*, vol. 6, no. 2, September 2006.
- [9] N. Patton, T. M. Aslam, T. MacGillivray, I. J. Deary, B. Dhillon, R. H. Eikelboom, K. Yogesan, and I. J. Constable, “Retinal image analysis: Concepts, applications and potential,” *Progress in Retinal and Eye Research*, vol. 25, pp. 99–127, 2006.
- [10] S. Hadjidemetriou, D. Toomre, and J. Duncan, “Microtubule segmentation and 3d reconstruction in total internal reflection fluorescence microscopy (tirfm),” in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2005, pp. 761–769.

- [11] H. Li, T. Shen, D. Vavylonis, and X. Huang, "Actin filament segmentation using spatiotemporal active-surface and active-contour models," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2010, pp. 86–94.
- [12] E. Meijering, "Neuron tracing in perspective," *Cytometry Part A*, vol. 77, no. 7, pp. 693–704, July 2010.
- [13] K. Brown, G. Barrionuevo, A. Canty, V. De Paola, J. Hirsch, G. Jefferis, J. Lu, M. Snippe, I. Sugihara, and G. Ascoli, "The diadem data sets: Representative light microscopy images of neuronal morphology to advance automation of digital reconstructions," *Neuroinformatics*, 2011.
- [14] C. Steger, "An unbiased detector of curvilinear structures," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, pp. 113–125, 1998.
- [15] F. Tupin, H. Maître, J.-F. Mangin, J.-M. Nicolas, and E. Pecherskyn, "Detection of linear features in sar images: Application to road network extraction," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, no. 2, pp. 434–453, March 1998.
- [16] G. Mountrakis, J. Im, and C. Ogole, "Support vector machines in remote sensing: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 66, no. 3, pp. 247–259, 2011.
- [17] N. Komodakis, N. Paragios, and G. Tziritas, "Clustering via lp-based stabilities," in *Neural Information Processing Systems (NIPS)*, 2008.
- [18] P. Torr and A. Zisserman, "Robust computation and parametrization of multiple view relations," in *International Conference on Computer Vision (ICCV)*, 1998, pp. 727–732.
- [19] C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer, "Optimizing binary mrfs via extended roof duality," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [20] A. A. Amini, T. E. Weymouth, and R. Jain, "Using dynamic programming for solving variational problems in vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, 1990.
- [21] T. Hauke Heibel, B. Glocker, M. Groher, N. Paragios, N. Komodakis, and N. Navab, "Discrete tracking of parametrized curves," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 1754–1761.
- [22] 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*, vol. 54, pp. 446–450, 2009.

- [23] P. Wang, W.-s. Liao, T. Chen, S. Kevin Zhou, and D. Comaniciu, "Graph based interactive detection of curve structures in 2d fluoroscopy," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2010, pp. 269–277.
- [24] A. Barbu, V. Athitsos, B. Georgescu, S. Boehm, P. Durlak, and D. Comaniciu, "Hierarchical learning of curves application to guidewire localization in fluoroscopy," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [25] G. Strang, *Linear algebra and its applications.*, 3rd ed. San Diego, CA: Harcourt Brace Jovanovich, 1988.
- [26] S. A. M. Baert, M. A. Viergever, and W. J. Niessen, "Guide-wire tracking during endovascular interventions," *IEEE Transactions on Medical Imaging*, vol. 22, no. 8, pp. 965–972, aug. 2003.
- [27] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, "Multiscale vessel enhancement filtering," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 1998.
- [28] P. Kovesi, "Image features from phase congruency," *Videre: A Journal of Computer Vision Research. MIT Press*, vol. 1, no. 3, 1999.
- [29] G. Slabaugh, K. Kong, G. Unal, and T. Fang, "Variational guidewire tracking using phase congruency," in *Medical Image Computing and Computer Assisted Intervention (MICCAI), Part II*, ser. LNCS, vol. 4792. Springer, 2007, pp. 612–619.
- [30] Z. Tu, "Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering," in *International Conference on Computer Vision (ICCV)*, 2005.
- [31] P. Wang, T. Chen, Y. Zhu, W. Zhang, S. K. Zhou, and D. Comaniciu, "Robust guidewire tracking in fluoroscopy," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [32] T. Hauke Heibel, B. Glocker, N. Paragios, and N. Navab, "Needle tracking through higher-order mrf optimization," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2010.
- [33] O. Pauly, T. Hauke Heibel, and N. Navab, "A machine learning approach for deformable guide-wire tracking in fluoroscopic sequences," in *Medical Image Computing and Computer Assisted Intervention (MICCAI), Part III*, ser. LNCS, vol. 6363. Springer, 2010, pp. 343–350.
- [34] E. Candès, "Ridgelets: Theory and applications," Ph.D. dissertation, Yale University, 1998.

- [35] D. L. Donoho, "Orthonormal ridgelets and linear singularities," *SIAM Journal on Mathematical Analysis*, vol. 31, pp. 1062–1099, 1998.
- [36] D. G. Lowe, "Object recognition from local scale-invariant features," *Proceedings of the International Conference on Computer Vision*, vol. 2, pp. 1150–1157, 1999.
- [37] —, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [38] C. Harris and M. Stevens, "A combined corner and edge detector," in *4th Alvey Vision Conference*, 1988.
- [39] W. Cai and A. C. Chung, "Multi-resolution vessel segmentation using normalized cuts in retinal images," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2006.
- [40] D. Marr and E. Hildreth, "Theory of edge detection," *Proceedings of the Royal Society of London. Series B, Biological Sciences*, vol. 207, no. 1167, pp. 187–217, February 1980.
- [41] C. Lorenz, I.-C. Carlsen, T. M. Buzug, C. Fassnacht, and J. Wesse, "Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images," *Proceedings of the 1st Joint Conference CVRMed and MRCAS (CVRMed/MRCAS97)*, pp. 233–242, 1997.
- [42] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis, "3-d multi-scale line filter for segmentation and visualization of curvilinear structures in medical images," *Medical Image Analysis*, vol. 2, no. 2, pp. 143–168, 1998.
- [43] R. Polli and G. Valli, "An algorithm for real-time vessel enhancement and detection," *Computer Methods and Programs in Biomedicine*, vol. 52, no. 1, pp. 1–22, January 1997.
- [44] S. Bouix, K. Siddiqi, and A. Tannenbaum, "Flux driven automatic centerline extraction," *Medical Image Analysis*, vol. 9, no. 3, pp. 209–221, 2005.
- [45] A. Vasilevskiy and K. Siddiqi, "Flux maximizing geometric flows," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 24, no. 12, pp. 1565–1578, December 2002.
- [46] M. Law and A. Chung, "Efficient implementation for spherical flux computation and its application to vascular segmentation," *IEEE Transactions on Image Processing*, vol. 18, no. 3, pp. 596–612, 2009.

- [47] C. Xu and J. L. Prince, "Gradient vector flow: A new external force for snakes," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1997, pp. 66–71.
- [48] C. Bauer and H. Bischof, "A novel approach for detection of tubular objects and its application to medical image analysis." in *Pattern Recognition*, 2008, pp. 163–172.
- [49] Y. Wang, A. Narayanaswamy, C. Tsai, and B. Roysam, "A broadly applicable 3-d neuron tracing method based on open-curve snake," *Neuroinformatics*, vol. 9, no. 2-3, pp. 193–217, 2011.
- [50] C. Bauer and H. Bischof, "Extracting curve skeletons from gray value images for virtual endoscopy." in *Medical Imaging and Augmented Reality*, 2008, pp. 393–402.
- [51] J. Serra, *Image Analysis and Mathematical Morphology*, 1982.
- [52] 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)*, vol. 7259, Feb. 2009.
- [53] O. Tankyevych, H. Talbot, P. Dokl  dal, and N. Passat, "Spatially-variant morpho-hessian filter : efficient implementation and applications." *International Symposium on Mathematical Morphology*, vol. 5720, pp. 137–148, 2009.
- [54] O. Tankyevych, "Filtering of thin objects, application to vascular images analysis," Ph.D. dissertation, University Paris-Est, 2010.
- [55] H. Talbot and B. Appleton, "Efficient complete and incomplete path openings and closings," *Image and Vision Computing*, vol. 25, pp. 416–425, 2007.
- [56] D. L. Donoho and X. Huo, "Beamlets and multiscale image analysis," Department of Statistics, Stanford University, Tech. Rep., April 2001.
- [57] S. Carlsson, "Sketch based coding of grey level images," *Signal Processing*, vol. 15, 1998.
- [58] E. J. Cand  s and D. Donoho, "Curvelets and curvilinear integrals," Department of Statistics, Stanford University, Tech. Rep., 1999.
- [59] M. N. Do and M. Vetterli, "The contourlet transform: An efficient directional multiresolution image representation," *IEEE Transactions on Image Processing*, vol. 14, no. 12, pp. 2091–2106, December 2005.
- [60] A. L. da Cunha, J. Zhou, and M. N. Do, "The nonsubsampling contourlet transform: Theory, design, and applications," *IEEE Transactions on Image Processing*, vol. 15, no. 10, pp. 3089–3101, 2006.

- [61] V. Velisavljević, B. Beferull-Lozano, M. Vetterli, and P. Luigi Dragotti, "Directionlets: Anisotropic multidirectional representation with separable filtering," *IEEE Transactions on Image Processing*, vol. 15, no. 7, pp. 1916–1933, July 2006.
- [62] P. Feng, Y. Pan, B. Wei, W. Jin, and D. Mi, "Enhancing retinal image by the contourlet transform," *Pattern Recognition Letters*, vol. 28, pp. 516–522, 2007.
- [63] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, 1991.
- [64] M. Jacob and M. Unser, "Design of steerable filters for feature detection using canny-like criteria," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 1007–1019, 2004.
- [65] G. Gonzàlez, F. Fleuret, and P. Fua, "Learning rotational features for filament detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [66] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.
- [67] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson, and M. Goldbaum, "Detection of blood vessels in retinal images using two-dimensional matched filters," *IEEE Transactions on Medical Imaging*, vol. 3, pp. 263–269, September 1989.
- [68] S. Hadjidemetriou, J. Duncan, D. Toomre, and D. Tuck, "Automatic quantification of microtubule dynamics," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2004, pp. 656–659.
- [69] A. Can, H. Shen, J. Turner, H. Tanenbaum, and B. Roysam, "Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms," *IEEE Transactions on Information Technology in Biomedicine*, vol. 3, no. 2, pp. 125–138, june 1999.
- [70] K. Al-Kofahi, S. Lasek, D. Szarowski, C. Pace, G. Nagy, J. Turner, and B. Roysam, "Rapid automated three-dimensional tracing of neurons from confocal image stacks," *IEEE Transactions on Information Technology in Biomedicine*, vol. 6, no. 2, pp. 171–187, June 2002.
- [71] R. Toledo, X. Orriols, P. Radeva, X. Binefa, J. Vitrià, C. Morales, and J. Villanueva, "Eigensnakes for vessel segmentation in angiography," in *International Conference on Pattern Recognition (ICPR)*, 2000, pp. 340–343.

- [72] J. Weickert, *A review of nonlinear diffusion filtering*, B. ter Haar Romeny, L. Florack, J. Koenderink, and M. Viergever, Eds. Berlin: Springer, 1997, vol. 1252, scale-Space Theory in Computer Vision, Lecture Notes in Computer Science.
- [73] G. Guy and G. Medioni, "Inference of surfaces, 3d curves and junctions from sparse, noisy 3d data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 11, pp. 1265–1277, November 1997.
- [74] L. Alvarez, P.-L. Lions, and J.-M. Morel, "Image selective smoothing and edge detection by nonlinear diffusion. ii." *SIAM Journal on Numerical Analysis*, vol. 29, pp. 845–866, 1992.
- [75] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629–639, July 1990.
- [76] J. Weickert, "Anisotropic diffusion in image processing," Ph.D. dissertation, University of Kaiserslautern., 1996.
- [77] ———, "Theoretical foundations of anisotropic diffusion in image processing," *Theoretical Foundations of Computer Vision*, pp. 221–236, 1996.
- [78] ———, "Coherence-enhancing diffusion filtering," *International Journal of Computer Vision*, vol. 31, pp. 111–127, 1999.
- [79] H. Kramer and J. Bruckner, "Iterations of a non-linear transformation for enhancement of digital images." *Pattern Recognition*, vol. 7, pp. 53–58, 1975.
- [80] S. Osher and L. Rudin, "Feature-oriented image enhancement using shock filters," *SIAM Journal on Numerical Analysis*, vol. 27, pp. 919–940, 1990.
- [81] G. Gilboa, N. Sochen, and Y. Y. Zeevi, "Image enhancement and denoising by complex diffusion processes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 1020–1036, 2004.
- [82] E. Meijering, W. Niessen, J. Weickert, and M. Viergever, "Evaluation of diffusion techniques for improved vessel visualization and quantification in three-dimensional rotational angiography," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2001.
- [83] R. Manniesing, M. A. Viergever, and W. J. Niessen, "Vessel enhancing diffusion a scale space representation of vessel structures," *Medical Image Analysis*, vol. 10, pp. 815–825, 2006.
- [84] H. M. Luu, A. Moelker, C. Klink, A. Mendrik, W. Niessen, and T. van Walsum, "Evaluation of diffusion filters for 3d cta liver vessel enhancement," *Abdominal Imaging. Computational and Clinical Applications*, pp. 168–177, 2012.

- [85] G. Guy and G. Medioni, “Inferring global perceptual contours from local features,” *International Journal of Computer Vision*, vol. 20, pp. 113–133, 1996.
- [86] G. Medioni, M. Lee, and C. Tang, *A Computational Framework for Segmentation and Grouping*. Elsevier, 2000.
- [87] P. Mordohai, “A perceptual organization approach for figure completion, binocular and multiple-view stereo and machine learning using tensor voting,” Ph.D. dissertation, Faculty of the graduate School, University of Southern California, 2005.
- [88] P. Mordohai and G. Medioni, “Dimensionality estimation, manifold learning and function approximation using tensor voting,” *Journal of Machine Learning Research*, vol. 11, pp. 411–450, 2010.
- [89] E. Jurrus, S. Watanabe, R. J. Giuly, A. R. C. Paiva, M. H. Ellisman, E. M. Jorgensen, and T. Tasdizen, “Semi-automated neuron boundary detection and nonbranching process segmentation in electron microscopy images,” *Neuroinformatics*, March 2012.
- [90] M. Wertheimer, “Untersuchungen zur lehre von der gestalt. i. prinzipielle bemerkungen.” *Psychologische Forschung*, vol. 1, pp. 47–58, 1922.
- [91] ———, “Untersuchungen zur lehre von der gestalt. ii,” *Psychologische Forschung*, vol. 4, pp. 301–350, 1923.
- [92] G. Kanizsa, *Gramática de la Visión*. Paidós, 1986.
- [93] L. Williams and D. Jacobs, “Stochastic completion fields: a neural model of illusory contour shape and salience,” *Neural Computation*, vol. 9, pp. 849–870, 1997.
- [94] K. Thornber and L. Williams, “Analytic solution of stochastic completion field,” *BioCyber*, vol. 75, pp. 141–151, 1996.
- [95] M. van Almsick, R. Duits, E. Franken, and B. t. Haar Romeny, “From stochastic completion fields to tensor voting,” *International workshop Deep structure, singularities, and computer vision (DSSCV)*, pp. 124–134, 2005.
- [96] T.-P. Wu, J. Jia, and C.-K. Tang, “A closed-form solution to tensor voting for robust parameter estimation via expectation-maximization,” Hong Kong University of Science and Technology, Tech. Rep., 2009.
- [97] E. Franken, M. van Almsick, P. Rongen, L. Florack, and B. ter Haar Romeny, “An efficient method for tensor voting using steerable filters,” in *European Conference on Computer Vision (ECCV)*, 2006.

- [98] R. Moreno, M. A. Garcia, D. Puig, L. Pizarro, B. Burgeth, and J. Weickert, "On improving the efficiency of tensor voting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 11, pp. 2215–2228, November 2011.
- [99] Y. Rouchdy and L. D. Cohen, "A geodesic voting method for the segmentation of tubular tree and centerlines," in *IEEE International Symposium on Biomedical Imaging (ISBI'11)*, Chicago, USA, April 2011, pp. 979–983.
- [100] V. Bismuth, R. Vaillant, H. Talbot, and L. Najman, "Curvilinear structure enhancement with the polygonal path image - application to guide-wire segmentation in xray fluoroscopy," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2012.
- [101] V. Vapnik, S. E. Golowich, and A. Smola, "Support vector method for function approximation, regression estimation, and signal processing," in *Advances in Neural Information Processing Systems*, vol. 9, 1996.
- [102] J. V. B. Soares, J. J. G. Leandro, R. M. Cesar Jr., H. F. Jelinek, and M. J. Cree, "Retinal vessel segmentation using the 2-d gabor wavelet and supervised classification," *IEEE Transactions on Medical Imaging*, vol. 25, no. 9, pp. 1214–1222, 2006.
- [103] J. V. B. Soares, J. J. Leandro, R. M. Cesar Jr., H. F. Jelinek, and M. J. Cree, "Using the 2-d morlet wavelet with supervised classification for retinal vessel segmentation," in *the Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI)*, 2005.
- [104] G. Gonzàlez, F. Aguet, F. Fleuret, M. Unser, and P. Fua, "Steerable features for statistical 3d dendrite detection," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2009.
- [105] G. Gonzàlez, E. Türetken, F. Fleuret, and P. Fua, "Delineating trees in noisy 2d images and 3d image-stacks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [106] E. Türetken, G. Gonzàlez, C. Blum, and P. Fua, "Automated reconstruction of dendritic and axonal trees by global optimization with geometric priors," *Neuroinformatics*, vol. 9, pp. 279–302, 2011.
- [107] M. Aizerman, E. Braverman, and L. Rozonoer, "Theoretical foundations of the potential function method in pattern recognition learning," *Automation and Remote Control*, vol. 25, pp. 821–837, 1964.
- [108] B. E. Boser, I. M. Guyon, and V. N. Vapnik, "A training algorithm for optimal margin classifiers," in *Fifth Annual Workshop on Computational Learning Theory*, 1992, pp. 144–152.

- [109] Y. Freund and R. Shapire, “A decision-theoretic generalization of on-line learning and an application to boosting,” in *International Conference on Machine Learning (ICML)*, 1996.
- [110] J. Friedman, T. Hastie, and R. Tibshirani, “additive logistic regression : a statistical view of boosting,” *the annals of statistics*, vol. 28, no. 2, pp. 337–407, 1998.
- [111] Y. Freund, R. Iyer, R. E. Schapire, and Y. Singer, “An efficient boosting algorithm for combining preferences,” *Journal of Machine Learning Research*, vol. 4, pp. 933–969, 2003.
- [112] N. Duffy and D. Helmbold, “Boosting methods for regression,” *Machine Learning*, vol. 47, pp. 153–200, 2002.
- [113] R. Meir and G. Rätsch, “An introduction to boosting and leveraging,” *Advanced Lectures on Machine Learning*, vol. LNAI 2600, pp. 118–183, 2003.
- [114] P. Viola and M. Jones, “Rapid object detection using a boosted cascade of simple features,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001.
- [115] P. Dollár, Z. Tu, and S. Belongie, “Supervised learning of edges and object boundaries,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [116] L. Breiman and E. Schapire, “Random forests,” *Machine Learning*, pp. 5–32, 2001.
- [117] Y.-S. Shih, “Families of splitting criteria for classification trees,” *Statistics and Computing*, vol. 9, pp. 309–315, 1999.
- [118] R. Rigamonti and V. Lepetit, “Accurate and efficient linear structure segmentation by leveraging ad hoc features with learned filters,” in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2012.
- [119] M. W. Law and A. C. Chung, “Three dimensional curvilinear structure detection using optimally oriented flux,” in *European Conference on Computer Vision (ECCV), Part IV*, 2008, pp. 368–382.
- [120] T. Hastie and R. Tibshirani, *Generalized Additive Models*, Chapman and Hall, Eds., 1990.
- [121] D. Palti-Wasserman, A. M. Bruckstein, and R. P. Beyar, “Identifying and tracking a guide wire in the coronary arteries during angioplasty from x-ray images,” *IEEE Transactions on Biomedical Engineering*, vol. 44, no. 2, pp. 152–164, feb. 1997.

- [122] S. A. M. Baert, E. B. van de Kraats, T. van Walsum, M. A. Viergever, and W. J. Niessen, “Three-dimensional guide-wire reconstruction from biplane image sequences for integrated display in 3-d vasculature,” *IEEE Transactions on Medical Imaging*, vol. 22, no. 10, pp. 1252–1258, October 2003.
- [123] J. Wu, S. C. Brubaker, M. D. Mullin, and J. M. Rehg, “Fast asymmetric learning for cascade face detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 3, pp. 369–382, March 2008.
- [124] K.-K. Sung and T. Poggio, “Example-based learning for view-based human face detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 1, pp. 39–51, 1998.
- [125] G. M. Weiss, “Mining with rarity: A unifying framework,” *ACM SIGKDD Explorations Newsletter - Special issue on learning from imbalanced datasets*, vol. 6, no. 1, pp. 7–19, June 2004.
- [126] P. Viola and M. Jones, “Fast and robust classification using asymmetric adaboost and a detector cascade,” in *Neural Information Processing Systems (NIPS)*, 2001.
- [127] G. Karakoulas and J. Shawe-Taylor, “Optimizing classifiers for imbalanced training sets,” in *Neural Information Processing Systems (NIPS)*, 1998.
- [128] W. Fan, S. J. Stolfo, J. Zhang, and P. K. Chan, “Adacost: misclassification cost-sensitive boosting,” in *Proceedings of the Sixteenth International Conference on Machine Learning (ICML)*, 1999, pp. 99–105.
- [129] M. V. Joshi, V. Kumar, and R. C. Agarwal, “Evaluating boosting algorithms to classify rare cases: comparison and improvements,” in *First IEEE International Conference on Data Mining*, November 2001, pp. 257–264.
- [130] E. Meijering, M. Jacob, J.-C. Sarria, P. Steiner, H. Hirling, and M. Unser, “Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images,” *Cytometry*, vol. 58A, pp. 167–176, 2004.
- [131] E. Franken, P. Rongen, M. van Almsick, and B. M. ter Haar Romeny, “Detection of electrophysiology catheters in noisy fluoroscopy images,” in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2006, pp. 25–32.
- [132] N. Honnorat, R. Vaillant, and N. Paragios, “Robust guidewire segmentation through boosting, clustering and linear programming,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2010, pp. 924–927.
- [133] —, “Guide-wire extraction through perceptual organization of local segments in fluoroscopic images,” in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2010.

- [134] G. González, F. Fleuret, and P. Fua, “Automated delineation of dendritic networks in noisy image stacks,” in *European Conference on Computer Vision (ECCV)*, 2008.
- [135] J. Mac Cullagh, P. Nelder, *Generalized Linear Models*, Chapman and Hall, Eds., 1989.
- [136] Z. Leng, J. Korenberg, B. Roysam, and T. Tasdizen, “A rapid 2-d centerline extraction method based on tensor voting,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2011.
- [137] J. Staal, S. N. Kalitzin, and M. A. Viergever, “A trained spin-glass model for grouping of image primitives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 7, pp. 1172–1182, 2005.
- [138] M. Sofka and C. V. Stewart, “Retinal vessel centerline extraction using multiscale matched filters, confidence and edge measures,” *IEEE Transactions on Medical Imaging*, vol. 25, no. 12, pp. 1531–1546, december 2006.
- [139] N. Honnorat, R. Vaillant, and N. Paragios, “Graph-based guide-wire segmentation through fusion of contrast-enhanced and fluoroscopic images,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2012.
- [140] T. Lindeberg, “Edge detection and ridge detection with automatic scale selection,” *International Journal of Computer Vision*, vol. 30, pp. 117–156, November 1998.
- [141] K. Krissian, G. Malandain, R. Ayache, N. Vaillant, and Y. Troussset, “Model-based detection of tubular structures in 3d images,” *Computer Vision and Image Understanding*, vol. 80, no. 2, pp. 130–171, 2000.
- [142] T. Lindeberg, *Scale-Space theory in Computer Vision*. Dordrecht, Netherlands: Kluwer Academic Publishers, 1994.
- [143] T. M. Koller, G. gerig, G. Szekely, and D. Dettwiler, “Multiscale detection of curvilinear structures in 2d and 3d image data,” in *International Conference on Computer Vision (ICCV)*, 1995, pp. 864–869.
- [144] C. Lorenz, I. C. Carse, T. M. Buzug, C. Fassnacht, and J. Weese, “Multiscale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images,” in *Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Robotics and Computer-Assisted Surgery*, March 1997, pp. 213–222.
- [145] J. Suri, K. Liu, L. Reden, and S. Laxminarayan, “A review on mr vascular image processing: skeleton versus nonskeleton approaches: part ii.” *IEEE Transactions on Information Technology in Biomedicine*, vol. 6, no. 4, pp. 338–350, 2002.

- [146] M. Schaap, C. Metz, T. van Walsum, A. van der Giessen, A. Weustink, N. Mollet, C. Bauer, H. Bogunović, C. Castro, X. Deng, E. Dikici, T. ODonnell, M. Frenay, O. Friman, M. Hoyos, P. Kitslaar, K. Krissian, C. Khnel, M. Orkisz, M. Luengo-Oroz, O. Smedby, M. Styner, A. Szymczak, H. Tek, C. Wang, S. Warfield, S. Zambal, and Y. Zhang, “Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms,” *Medical Image Analysis*, vol. 13, no. 5, pp. 701–714, 2009.
- [147] J. Howard and A. Hyman, “Dynamics and mechanics of the microtubule plus end,” *Nature*, vol. 422, pp. 753–758, 2003.
- [148] D. Axelrod, “Biophysical tools for biologists, volume two: In vivo techniques ; chapter 7 : Total internal reflection fluorescence microscopy,” *Methods in Cell Biology*, vol. 89, pp. 169–221, 2008.
- [149] M. E. Sargin, A. Altinok, K. Rose, and B. Manjunath, “Tracing curvilinear structures in live cell images,” in *IEEE International Conference on Image Processing (ICIP)*, 2007.
- [150] I. Smal, K. Draegestein, N. Galjart, W. Niessen, and E. Meijering, “Particle filtering for multiple object tracking in dynamic fluorescence microscopy images: Application to microtubule growth analysis.” *IEEE Transactions on Medical Imaging*, vol. 27, pp. 789–804, 2008.
- [151] H. Li, D. Shen, Tian Vavylonis, and X. Huang, “Actin filament tracking based on particle filters and stretching open active contour models,” in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2009.
- [152] S. Reuman and J. Capowski, “Automated neuron tracing using the marr-hildreth zerocrossing technique,” *Computers and Biomedical Research*, vol. 17, no. 2, pp. 93–115, April 1984.
- [153] G. Hui Zhang and N. Adluru, in *CDMRI’11: MICCAI 2011 Workshop on Computational Diffusion MRI*, 2011.
- [154] S. Pujol, R. Kikinis, M. Styner, A. Golby, A. Nabavi, G. Gerig, W. Wells, C.-F. Westin, and S. Gouttard, in *MICCAI 2012 DTI Tractography Challenge*, 2012.
- [155] G. Ascoli, K. Svoboda, and Y. Liu. (2009-2010) The diadem challenge. <http://www.diademchallenge.org/>.
- [156] *Workshop CDMRI 2011, organised during MICCAI*. Hui Zhang, Gary and Adluru, Nagesh, 2011.
- [157] C. Steger, “An unbiased detector of curvilinear structures,” *Technisch Universität München, München, Germany, Tech. Rep.*, July 1996.

- [158] R. Collorec and J. Coatrieux, "Vectorial tracking and directed contour finder for vascular network in digital subtraction angiography." *Pattern Recognition Letters*, vol. 8, no. 5, pp. 353–358, 1988.
- [159] J. Coatrieux, M. Garreau, R. Collorec, and C. Roux, "Computer vision approaches for the three-dimensional reconstruction of coronary arteries: review and prospects." *Critical Reviews in Biomedical Engineering*, vol. 22, no. 1, pp. 1–38, 1994.
- [160] K. Al-Kofahi, A. Can, S. Lasek, D. Szarowski, N. Dowell-Mesfin, W. Shain, J. Turner, and B. Roysam, "Median-based robust algorithms for tracing neurons from noisy confocal microscope images," *IEEE Transactions on Information Technology in Biomedicine*, vol. 7, no. 4, pp. 302–317, December 2003.
- [161] Y. Zhang, X. Zhou, A. Degterev, M. Lipinski, D. Adjero, J. Yuan, and S. T. Wong, "A novel tracing algorithm for high throughput imaging screening of neuron-based assays," *Journal of Neuroscience Methods*, vol. 160, pp. 149–162, 2007.
- [162] S. Aylward and E. Bullitt, "Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction," *IEEE Transactions on Medical Imaging*, vol. 21, no. 2, pp. 61–75, 2002.
- [163] N. Flaque, M. Desvignes, J. Constans, and M. Revenu, "Acquisition, segmentation and tracking of the cerebral vascular tree on 3d magnetic resonance angiography images." *Medical Image Analysis*, vol. 5, no. 3, pp. 173–183, 2001.
- [164] P. Reuzé, J. Coatrieux, L. Luo, and J. Dillenseger, "A 3-d moment based approach for blood vessel detection and quantification in mra," INRIA, France, Tech. Rep., 1993.
- [165] C. Bołdak, Y. Rolland, and C. Toumoulin, "An improved model-based vessel tracking algorithm with application to computed tomography angiography," *Biocybernetics and Biomedical Engineering*, vol. 23, no. 1, pp. 41–63, 2003.
- [166] A. Larralde, C. Bołdak, M. Garreau, C. Toumoulin, D. Boulmier, and Y. Rolland, "Evaluation of a 3d segmentation software for the coronary characterization in multi-slice computed tomography," in *Functional Imaging and Modeling of the Heart*, 2003, pp. 39–51.
- [167] T. Hauke-Heibel, B. Glocker, M. Groher, M. Pfister, and N. Navab, "Interventional tool tracking using discrete optimization," *IEEE Transactions on Medical Imaging*, December 2012.

- [168] J. Tyrrell, E. di Tomaso, D. Fuja, T. Ricky, K. Kozak, R. Jain, and B. Roysam, "Robust 3-d modeling of vasculature imagery using superellipsoids," *IEEE Transactions on Medical Imaging*, vol. 26, no. 2, pp. 223–237, February 2007.
- [169] A. H. Barr, "Superquadrics and angle-preserving transformations," *IEEE Computer Graphics and Applications*, vol. 1, no. 1, pp. 11–23, January 1981.
- [170] P. Fua and Y. Leclerc, "Model driven edge detection." *Machine Vision and Applications*, vol. 3, pp. 45–56, 1990.
- [171] J. Shi and J. Malik, "Normalized cuts and image segmentation." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, pp. 888–905, 2000.
- [172] D. Reid, "An algorithm for tracking multiple targets." *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.
- [173] O. Friman, M. Hindennach, C. Khnel, and H.-O. Peitgen, "Multiple hypothesis template tracking of small 3d vessel structures," *Medical Image Analysis*, vol. 14, pp. 160–171, 2010.
- [174] L. Liang, H. Shen, P. De Camilli, and J. S. Duncan, "Tracking clathrin coated pits with a multiple hypothesis based method." in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2010, pp. 315–322.
- [175] S. Blackman, "Multiple hypothesis tracking for multiple target tracking," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 5–18, 2004.
- [176] O. Friman, M. Hindennach, and H.-O. Peitgen, "Template-based multiple hypotheses tracking of small vessels," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2008, pp. 1047–1050.
- [177] R. Kalman, "A new approach to linear filtering and prediction problems," *ASME Journal of Basic Engineering*, vol. 82, no. D, pp. 35–45, 1960.
- [178] S. Wörz and K. Rohr, "A new 3d parametric intensity model for accurate segmentation and quantification of human vessels," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2004, pp. 491–499.
- [179] —, "Segmentation and quantification of human vessels using a 3-d cylindrical intensity model," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 1994–2004, August 2007.
- [180] D. Terzopoulos and R. Szeliski, "Tracking with kalman snakes," *Active Vision*, MIT Press, Cambridge, MA, pp. 3–20, 1992.

- [181] N. Peterfreund, "Robust tracking of position and velocity with kalman snakes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 6, pp. 564–569, 1999.
- [182] E. Jurrus, M. Hardy, T. Tasdizen, P. T. Fletcher, P. Koshevoy, C.-B. Chien, W. Denk, and R. Whitaker, "Axon tracking in serial block-face scanning electron microscopy," *Medical Image Analysis*, vol. 13, no. 1, pp. 180–188, February 2009.
- [183] C. Gössl, L. Fahrmeir, B. Pütz, L. Auer, and D. Auer, "Fibertracking from dti using linearstatespace models: Detectability of the pyramidaltract," *NeuroImage*, vol. 16, no. 2, pp. 378–388, 2002.
- [184] R. Brown and P. Hwang, *Introduction to random signals and applied Kalman Filtering*, second edition ed. Wiley, 1992.
- [185] G. Welch and G. Bishop, "An introduction to the kalman filter." Department of Computer Science, University of North Carolina - Chapel Hill, TR95-041., Tech. Rep., 1995.
- [186] J. Zhou, W. Bischof, and T. Caelli, "Road tracking in aerial images based on human-computer interaction and bayesian filtering," *ISPRS journal of photogrammetry and remote sensing*, vol. 61, pp. 108–124, 2006.
- [187] S. J. Julier and J. K. Uhlmann, "A new extension of the kalman filter to nonlinear systems," in *AeroSense: The 11th International Symposium on Aerospace/Defence Sensing, Simulation and Controls*, 1997.
- [188] E. Wan, R. van derMerwe, and A. T. Nelson, "Dual estimation and the unscented transformation," in *Neural Information Processing Systems (NIPS)*, M. Press, Ed., no. 12, 2000, pp. 666–672.
- [189] R. van der Merwe and E. Wan, "Sigma-point kalman filters for probabilistic inference in dynamic state-space models." in *Workshop on Advances in Machine Learning*, 2003.
- [190] S. Julier and J. Uhlmann, "Unscented filtering and nonlinear estimation," in *Proceedings of the IEEE*, no. 3, 2004, pp. 401–422.
- [191] J. G. Malcolm, M. E. Shenton, and Y. Rathi, "Neural tractography using an unscented kalman filter," in *Information Processing in Medical Imaging (IPMI)*, 2009.
- [192] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, February 2002.
- [193] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice.*, Springer-Verlag, Ed., 2001.

- [194] M. Isard and A. Blake, "Condensation conditional density propagation for visual tracking," *International Journal of Computer Vision*, vol. 29, 1998.
- [195] H. E. Çetingül, G. Plank, N. Trayanova, and R. Vidal, "Stochastic tractography in 3-d images via nonlinear filtering and spherical clustering," *MICCAI Workshop on Probabilistic Models for Medical, PMMIA*, pp. 268–279, 2009.
- [196] C. Florin, N. Paragios, and J. Williams, "A quasi-monte carlo solution for segmentation of coronaries." in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2005, pp. 246–253.
- [197] H. Shim, D. Kwon, I. Yun, and S. Lee, "Robust segmentation of cerebral arterial segments by a sequential monte carlo method: particle filtering." *Comput Methods and Programs in Biomedicine*, vol. 84, no. 2-3, pp. 135–145, 2006.
- [198] M. Schaap, I. Smal, C. Metz, T. van Walsum, and W. Niessen, "Bayesian tracking of elongated structures in 3d images," in *Information Processing in Medical Imaging (IPMI)*, 2007.
- [199] I. Smal, E. Meijering, K. Draegestein, N. Galjart, I. Grigoriev, A. Akhmanova, M. E. van Royen, A. B. Houtsmuller, and W. Niessen, "Multiple object tracking in molecular bioimaging by rao-blackwellized marginal particle filtering," *Medical Image Analysis*, vol. 12, no. 6, pp. 764–777, December 2008.
- [200] D. Lesage, E. D. Angelini, I. Bloch, and G. Funka-Lea, "Medial-based bayesian tracking for vascular segmentation: Application to coronary arteries in 3d ct angiography," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2008.
- [201] —, "Bayesian maximal paths for coronary artery segmentation from 3d ct angiograms," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2009.
- [202] M. Fischler, J. Tenenbaum, and H. Wolf, "Detection of roads and linear structures in low-resolution aerial imagery using a multisource knowledge integration technique," *Computer Graphics and Image Processing*, vol. 15, pp. 201–223, 1981.
- [203] H. Li, T. Shen, and X. Huang, "Actin filament segmentation using dynamic programming," in *Information Processing in Medical Imaging (IPMI)*, 2011.
- [204] P. F. Felzenszwalb and D. P. Huttenlocher, "Distance transforms of sampled functions." Cornell Computing and Information Science, Tech. Rep., 2004.
- [205] —, "Pictorial structures for object recognition," *International Journal of Computer Vision*, vol. 61, no. 1, pp. 55–79, 2005.

- [206] N. Komodakis and G. Tziritas, "Approximate labeling via graph-cuts based on linear programming," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2007.
- [207] N. Komodakis, G. Tziritas, and N. Paragios, "Performance vs computational efficiency for optimizing single and dynamic mrfs: Setting the state of the art with primal-dual strategies," *Computer Vision and Image Understanding*, vol. 112, pp. 14–29, 2008.
- [208] L. Cohen, "On active contour models and balloons," *CVGIP: Image understanding*, vol. 53, no. 2, pp. 211–218, 1991.
- [209] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models." *International Journal of Computer Vision*, vol. 1, pp. 321–331, 1988.
- [210] A. F. Frangi, W. J. Niessen, T. Hoogeveen, R. M. van Walsum, and M. A. Viergever, "Quantitation of vessel morphology from 3d mra," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 1999.
- [211] A. Frangi, W. Niessen, P. Nederkoorn, O. Elgersma, and M. Viergever, "Three-dimensional model-based stenosis quantification of the carotid arteries from contrast-enhanced mr angiography," in *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, 2000.
- [212] P. J. Yim, J. J. Cebal, R. Mullick, H. B. Marcos, and P. L. Choyke, "Vessel surface reconstruction with a tubular deformable model," *IEEE Transactions on Medical Imaging*, vol. 20, no. 12, pp. 1411–1421, December 2001.
- [213] J. Mille, R. Boné, and L. Cohen, "Region-based 2d deformable generalized cylinder for narrow structures segmentation." in *European Conference on Computer Vision (ECCV)*, Marseille, France, 2008, pp. 392–404.
- [214] F. Benmansour, L. Cohen, M. Law, and A. Chung, "Tubular anisotropy for 2d vessels segmentation," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Miami, Florida, USA, 2009.
- [215] I. Laptev, H. Mayer, T. Lindeberg, W. Eckstein, C. Steger, and A. Baumgartner, "Automatic extraction of roads from aerial images based on scale space and snakes," *Machine Vision and Applications*, vol. 12, no. 1, pp. 23–31, 2000.
- [216] S. C. Zhu and A. Yuille, "Region competition: Unifying snakes, region growing and bayes mdl for multiband image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 884–900, September 1996.
- [217] T. McInerney and D. Terzopoulos, "T-snakes: topology adaptive snakes," *Medical Image Analysis*, vol. 4, no. 2, pp. 73–91, 2000.

- [218] A. Dervieux and F. Thomasset, "A finite element method for the simulation of a rayleigh-taylor instability," *Lecture Notes in Mathematics*, vol. 771, pp. 145–158, 1980.
- [219] S. Osher and J. A. Sethian, "Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulation," *Journal of Computational Physics*, vol. 79, pp. 12–49, 1988.
- [220] T. Chan and L. Vese, "A multiphase level set framework for image segmentation using the mumford and shah model," *International Journal of Computer Vision*, vol. 50, no. 3, pp. 271–293, 2002.
- [221] L. Lorigo, O. Faugeras, W. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Westin, "Curves: curve evolution for vessel segmentation," *Medical Image Analysis*, vol. 5, pp. 195–206, 2001.
- [222] V. Luboz, X. Wu, K. Krissian, C.-F. Westin, R. Kikinis, S. Cotin, and S. Dawson, "A segmentation and reconstruction technique for 3d vascular structures," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2005.
- [223] M. Rochery, I. Jermyn, and J. Zerubia, "Higher order active contours," *International Journal of Computer Vision*, vol. 69, no. 1, pp. 27–42, 2006.
- [224] —, "Phase field models and higher-order active contours." in *IEEE International Conference on Computer Vision (ICCV)*, no. 2, 2005, pp. 970–976.
- [225] M. Descoteaux, L. Collins, and K. Siddiqi, "A multi-scale geometric flow for segmenting vasculature in mri," *Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis*, vol. 3117, pp. 169–180, 2004.
- [226] T. Chan and L. Vese, "Active contours without edges," *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 266–277, 2001.
- [227] R. Kimmel and A. M. Bruckstein, "Regularized laplacian zero crossings as optimal edge integrators," *International Journal of Computer Vision*, vol. 53, no. 3, pp. 225–243, 2003.
- [228] M. Holtzman-Gazit, R. Kimmel, N. Peled, and D. Goldsher, "Segmentation of thin structures in volumetric medical images," *IEEE Transactions on Image Processing*, vol. 15, no. 2, pp. 354–363, 2006.
- [229] D. Nain, A. Yezzi, and G. Turk, "Vessel segmentation using a shape driven flow," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2004.
- [230] P. Yan and A. A. Kassim, "Mra image segmentation with capillary active contour," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2005.

- [231] E. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, 1959.
- [232] M. L. Fredman and R. E. Tarjan, "Fibonacci heaps and their uses in improved network optimization algorithms," *Journal of the Association for Computing Machinery*, vol. 34, no. 3, pp. 596–615, July 1987.
- [233] J. Xie, T. Zhao, T. Lee, E. Myers, and H. Peng, "Automatic neuron tracing in volumetric microscopy images with anisotropic path searching." in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2010, pp. 472–479.
- [234] D. Johnson, "Efficient algorithms for shortest paths in sparse networks." *Journal of the Association for Computing Machinery*, vol. 1, pp. 1–13, 1977.
- [235] S. D. Olabarriaga, M. Breeuwer, and W. J. Niessen, "Minimum cost path algorithm for coronary artery central axis tracking in ct images." in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2003, pp. 687–694.
- [236] E. Jurrus, R. Whitaker, B. W. Jones, R. Marc, and T. Tasdizen, "An optimal-path approach for neural circuit reconstruction," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2008.
- [237] S. W. Zucker, "Region growing: Childhood and adolescence," *Computer Graphics and Image Processing*, vol. 5, no. 3, pp. 382–399, 1976.
- [238] R. Adams and L. Bischof, "Seeded region growing." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, no. 6, pp. 641–647, 1994.
- [239] J. Yi and J. Ra, "A locally adaptive region growing algorithm for vascular segmentation." *International Journal of Imaging Systems and Technology*, vol. 13, no. 4, pp. 208–214, 2003.
- [240] C. Metz, M. Schaap, A. G. van der Giessen, T. van Walsum, and W. J. Niessen, "Semi-automatic coronary artery centerline extraction in computed tomography angiography data." in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2007, pp. 856–859.
- [241] Y. Masutani, T. Schiemann, and K. Höhne, "Vascular shape segmentation and structure extraction using a shape-based region-growing model," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 1998.
- [242] S. Wan, A. Kiraly, E. Ritman, and W. Higgins, "Extraction of the hepatic vasculature in rats using 3-d micro-ct images." *IEEE Transactions on Medical Imaging*, vol. 19, no. 9, pp. 964–971, 2000.

- [243] S. Wan, E. Ritman, and W. Higgins, "Multi-generational analysis and visualization of the vascular tree in 3d micro-ct images." *Computers in Biology and Medicine*, vol. 32, no. 2, pp. 55–71, 2002.
- [244] T. Lee, R. Kashyap, and C. Chu, "Building skeleton models via (3-d) medial surface/axis thinning algorithms," *CVGIP: Graph. Models Image Process.*, vol. 56, pp. 462–478, 1994.
- [245] K. Siddiqui and S. M. Pizer, *Medial Representations : Mathematics, Algorithms and Applications*. Springer, 2008.
- [246] P. Saha, B. Chaudhuri, and D. D. Majumder, "A new shape preserving parallel thinning algorithm for 3d digital images," *Pattern Recognition*, vol. 30, no. 12, pp. 1939–1955, 1997.
- [247] K. Buhler, P. Felkel, and A. Cruz, "Geometric methods for vessel visualization and quantification - a survey." in *Geometric Modelling for Scientific Visualization*, Springer-Verlag, Ed., 2002.
- [248] N. Cornea, D. Silver, and P. Min, "Curve-skeleton properties, applications and algorithms," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 3, pp. 530–548, May-June 2007.
- [249] C. Kirbas and F. Quek, "Vessel extraction in medical images by 3d wave propagation and traceback." in *Third IEEE Symposium on Bioinformatics and Bioengineering, 2003. Proceedings.*, 2003, pp. 174–181.
- [250] O. Wink, W. Niessen, B. Verdonck, and M. Viergever, "Vessel axis determination using wave front propagation analysis." in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2001, pp. 845–853.
- [251] O. Wink, W. Niessen, and M. Viergever, "Multiscale vessel tracking." *IEEE Transactions on Medical Imaging*, vol. 23, no. 1, pp. 130–133, 2004.
- [252] R. Manniesing, M. A. Viergever, and W. J. Niessen, "Vessel axis tracking using topology constrained surface evolution," *IEEE Transactions on Medical Imaging*, vol. 26, no. 3, pp. 309–316, 2007.
- [253] J. N. Tsitsiklis, "Efficient algorithms for globally optimal trajectories," *IEEE Transactions on Automatic Control*, vol. 40, no. 9, pp. 1528–1538, September 1995.
- [254] D. Adalsteinsson and J. Sethian, "A fast level set method for propagating interfaces." *Journal of Computational Physics*, vol. 118, no. 2, pp. 269–277, 1995.
- [255] J. Sethian, *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry*, 2nd ed. Cambridge University Press, 1999.

- [256] ———, “A fast marching level set method for monotonically advancing fronts.” *Proceedings of the National Academy of Science*, vol. 93, no. 4, pp. 1591–1595, 1996.
- [257] T. Deschamps and L. D. Cohen, “Fast extraction of minimal paths in 3d images and applications to virtual endoscopy (2001),” *Medical Image Analysis*, vol. 5, p. 281299, 2001.
- [258] R. Manniesing and W. Niessen, “Local speed functions in level set based vessel segmentation.” in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2004, pp. 475–482.
- [259] L. Yatziv, A. Bartesaghi, and G. Sapiro, “O(n) implementation of the fast marching algorithm,” *Journal of Computational Physics*, vol. 212, pp. 393–399, 2006.
- [260] J. Yan and T. Zhuang, “Applying improved fast marching method to endocardial boundary detection in echocardiographic images.” *Pattern Recognition Letters*, vol. 24, no. 15, pp. 2777–2784, 2003.
- [261] T. Deschamps and L. D. Cohen, “Fast extraction of tubular and tree 3d surfaces with front propagation methods,” in *International Conference on Pattern Recognition (ICPR)*, 2002, pp. 731–734.
- [262] Q. Lin, “Enhancement, detection, and visualization of 3d volume data. dissertations no. 824,” Ph.D. dissertation, Dept. EE, Linkping University, SE-581 83 Linkping, Sweden, 2003.
- [263] F. Benmansour and L. D. Cohen, “Tubular structure segmentation based on minimal path method and anisotropic enhancement.” *International Journal of Computer Vision*, vol. 92, no. 2, pp. 192–210, April 2011.
- [264] Y. Zhu, X. Papademetris, A. J. Sinusas, and J. S. Duncan, “Local shape registration using boundary-constrained match of skeletons.” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [265] Y. Zhou, A. Kaufman, and A. Toga, “Three-dimensional skeleton and centerline generation based on an approximate minimum distance field.” *Visual Computer*, vol. 14, pp. 303–314, 1998.
- [266] Y. Zhou and A. Toga, “Efficient skeletonization of volumetric objects.” *IEEE Transactions on Visualization and Computer Graphics*, vol. 5, pp. 196–209, 1999.
- [267] Z. Vasilkoski and A. Stepanyants, “Detection of the optimal neuron traces in confocal microscopy images.” *Journal of Neuroscience Methods*, vol. 178, pp. 197–204, 2009.

- [268] H. Li and A. Yezzi, “Vessels as 4-d curves: global minimal 4-d paths to extract 3-d tubular surfaces and centerlines.” *IEEE Transactions on Medical Imaging*, vol. 26, no. 9, pp. 1213–1223, 2007.
- [269] M. Nitzberg, D. Mumford, and T. Shiota, “Filtering, segmentation and depth,” *Lecture Notes in Computer Science*, vol. 662, 1993.
- [270] D. Mumford, *Elastica in Computer Science*. Springer Verlag, 1994.
- [271] E. Sharon, A. Brandt, and R. Basri, “Completion energies and scale,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1117–1131, 2000.
- [272] I. Jermyn and H. Ishikawa, “Globally optimal regions and boundaries as minimum ratio weight cycles.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 10, pp. 1075–1088, 2001.
- [273] S. Wang, T. Kubota, J. M. Siskind, and J. Wang, “Salient closed boundary extraction with ratio contour,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 4, no. 27, pp. 546–561, 2005.
- [274] I. Kokkinos, “Highly accurate boundary detection and grouping,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [275] P. Felzenszwalb and D. McAllester, “A min-cover approach for finding salient curves,” in *Conference on Computer Vision and Pattern Recognition Workshop (CVPRW '06)*, 2006.
- [276] S. Z. Li, *Markov Random Field Modeling in Image Analysis*. Springer, 2009.
- [277] M. Negri, P. Gamba, G. Lisini, and F. Tupin, “Junction-aware extraction and regularization of urban road networks in high-resolution sar images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 10, pp. 2962–2971, 2006.
- [278] J. Staal, S. N. Kalitzin, and M. A. Viergever, “A trained spin-glass model for grouping of image primitives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 7, pp. 1172–1182, July 2005.
- [279] V. Murino, C. Regazzoni, and G. Foresti, “Grouping as a searching process for minimum-energy configurations of labelled random fields,” *Computer Vision and Image Understanding*, vol. 64, no. 1, pp. 157–174, 1996.
- [280] J. Lafferty, A. McCallum, and F. Pereira, “Conditional random fields: Probabilistic models for segmenting and labeling sequence data,” in *Eighteenth International Conference on Machine Learning (ICML01)*, 2001.
- [281] X. Ren, C. C. Fowlkes, and J. Malik, “Scale-invariant contour completion using conditional random fields,” in 2005,” in *IEEE International Conference on Computer Vision (ICCV)*, 2005.

- [282] R. Stoica, X. Descombes, and J. Zerubia, “A gibbs point process for road extraction from remotely sensed images,” *International Journal of Computer Vision*, vol. 57, no. 2, pp. 121–136, 2004.
- [283] C. Lacoste, G. Finet, and I. Magnin, “Coronary tree extraction from x-ray angiograms using marked point processes.” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2006.
- [284] P. Green, “Reversible jump mcmc computation and bayesian model determination.” *Biometrika*, vol. 82, pp. 711–732, 1995.
- [285] Z. Tu and S.-C. Zhu, “Parsing images into regions, curves, and curve groups,” *International Journal of Computer Vision*, vol. 69, no. 2, pp. 223–249, 2006.
- [286] L. Yatziv, M. Chartouni, S. Datta, and G. Sapiro, “Toward multiple catheters detection in fluoroscopic image guided interventions,” *IEEE Transactions on Information Technology in Biomedicine*, vol. 16, no. 4, pp. 770–781, July 2012.
- [287] A. DeLong, A. Osokin, H. N. Isack, and Y. Boykov, “Fast approximate energy minimization with label costs.” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [288] L. Ladicky, C. Russell, P. Kohli, and T. P. H.S., “Graph cut based inference with co-occurrence statistics,” in *European Conference on Computer Vision (ECCV)*, 2010.
- [289] E. Türetken, F. Benmansour, and P. Fua, “Automated reconstruction of tree structures using path classifiers and mixed integer programming,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [290] J. B. MacQueen, “Some methods for classification and analysis of multivariate observations.” in *Proceedings of the Fifth Symposium on Math, Statistics and Probability*, Berkeley, CA: University of California Press, 1967, pp. 281–297.
- [291] H. Akaike, “A new look at the statistical model identification,” *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716–723, 1974.
- [292] G. E. Schwarz, “Estimating the dimension of a model,” *Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.
- [293] N. Honnorat, R. Vaillant, J. S. Duncan, and N. Paragios, “Curvilinear structures extraction in cluttered bioimaging data with discrete optimization methods,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2011.

- [294] S. Pizer, C. Burbeck, J. Coggins, D. Fritsch, and B. Morse, “Object shape before boundary shape: Scalespace medial axes,” *Journal of Mathematical Imaging and Vision*, vol. 4, pp. 303–313, 1994.
- [295] T. Pock, C. Janko, R. Beichel, and H. Bischof, “Multiscale medialness for robust segmentation of 3d tubular structures,” in *in Proceedings of the Computer Vision Winter Workshop*, 2005, pp. 93–102.
- [296] H. Blum, *A transformation for extracting new descriptors of shape*. Cambridge: MIT Press, 1967, pp. 363–380.
- [297] H. Blum and R. Nagel, “Shape description using weighted symmetric axis features,” *Pattern Recognition*, vol. 10, pp. 167–180, 1978.
- [298] M. A. Fischler and R. C. Bolles, “Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography,” *Readings in computer vision: issues, problems, principles, and paradigms*, pp. 726–740, 1987.
- [299] P. Meer, *Emerging Topics in Computer Vision*. Prentice Hall, 2004, ch. 4 : Robust Techniques for Computer Vision.
- [300] V. Kolmogorov, “Convergent tree-reweighted message passing for energy minimization,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 10, pp. 1568–1583, October 2006.
- [301] M. Wainwright, T. Jaakkola, and A. Willsky, “Map estimation via agreement on trees: Message-passing and linear programming,” *IEEE Transactions on Information Theory*, vol. 51, no. 11, pp. 3697–3717, 2005.
- [302] A. Niculescu-Mizil and R. Caruana, “obtaining calibrated probabilities from boosting,” in *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence*, 2005.
- [303] H. Ishikawa, “Higher-order clique reduction in binary graph cut,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 2993–3000.
- [304] N. Komodakis, N. Paragios, and G. Tziritas, “Mrf optimization via dual decomposition: Message-passing revisited,” in *International Conference on Computer Vision (ICCV)*, 2007.
- [305] P. Pérez, A. Blake, and M. Gangnet, “Jetstream: probabilistic contour extraction with particles,” in *IEEE International Conference on Computer Vision (ICCV)*, 2001.
- [306] B. D. Lucas and T. Kanade, “An iterative image registration technique with an application to stereo vision,” *Proceedings of Imaging Understanding Workshop*, pp. 121–130, 1981.

- [307] M. G. Danilouchkine, F. Mastik, and A. F. W. van der Steen, "A study of coronary artery rotational motion with dense scale-space optical flow in intravascular ultrasound," *Physics in Medicine and Biology*, vol. 54, no. 6, p. 1397, 2009.
- [308] A. Hamadeh, S. Lavallé, and P. Cinquin, "Automated 3-dimensional computed tomographic and fluoroscopic image registration." *Computer Aided Surgery*, vol. 3, no. 1, 1998.
- [309] A. Guézic, P. Kazanzides, B. Williamson, and R. H. Taylor, "Anatomy-based registration of ct-scan and introoperative x-ray images for guiding surgical robot." *IEEE Transactions on Medical Imaging*, vol. 17, pp. 715–728, 1998.
- [310] M. Sargin, A. Altinok, K. Rose, and B. Manjunath, "Deformable trellis: Open contour tracking in bio-image sequences," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, April 2008.
- [311] A. Altinok, M. El-Saban, A. J. Peck, L. Wilson, S. C. Feinstein, B. S. Manjunath, and K. Rose, "Activity analysis in microtubule videos by mixture of hidden markov models," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [312] J. Shi and C. Tomasi, "Good features to track," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1994.
- [313] C. Tomasi and T. Kanade, "Detection and tracking of point features," *Pattern Recognition*, vol. 37, pp. 165–168, 2004.
- [314] C. Schmid, R. Mohr, and C. Bauckhage, "Evaluation of interest point detectors," *International Journal of Computer Vision*, vol. 37, no. 2, pp. 151–172, 2000.
- [315] F. Mokhtarian and F. Mohanna, "Performance evaluation of corner detectors using consistency and accuracy measures," *Computer Vision and Image Understanding*, vol. 102, no. 1, pp. 81–94, 2006.
- [316] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schafalitzky, T. Kadir, and L. Van Gool, "A comparison of affine region detectors," *International Journal of Computer Vision*, vol. 65, no. 1-2, pp. 43–72, 2005.
- [317] L. Juan and O. Gwun, "A comparison of sift , pca-sift and surf," *International Journal of Image Processing (IJIP)*, vol. 3, no. 5, pp. 143–152, 2010.
- [318] J. Heikkilä, "Pattern matching with affine moment descriptors," *Pattern Recognition*, vol. 37, no. 9, pp. 1825–1834, 2004.

- [319] J. Glaunes, A. Trouve, and L. Younes, “Diffeomorphic matching of distributions: a new approach for unlabelled point-sets and sub-manifolds matching,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004, pp. 712–718.
- [320] Y. Tsin and T. Kanade, “A correlation-based approach to robust point set registration,” in *European Conference on Computer Vision (ECCV)*, 2004, pp. 558–569.
- [321] M. Singh, A. Himanshu, and N. Ahuja, “Robust registration and tracking using kernel density correlation,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2004.
- [322] P. J. Besl and N. D. McKay, “A method for registration of 3-d shapes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [323] Z. Zhang, “Iterative point matching for registration of free-form curves and surfaces,” *International Journal of Computer Vision*, vol. 13, no. 2, pp. 119–152, 1994.
- [324] Y. Liu, “Improving icp with easy implementation for free-form surface matching,” *Pattern Recognition*, vol. 37, no. 2, pp. 211–226, February 2004.
- [325] H. Chui and A. Rangarajan, “A new point matching algorithm for non-rigid registration,” *Computer Vision and Image Understanding*, vol. 89, no. 2-3, pp. 114–141, March 2003.
- [326] M. Leordeanu and M. Hebert, “A spectral technique for correspondence problems using pairwise constraints,” in *IEEE International Conference on Computer Vision (ICCV)*, vol. 2, 2005, pp. 1482–1489.
- [327] I. C. Berg, T. L. Berg, and J. Malik, “Shape matching and object recognition using low distortion correspondences,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005, pp. 26–33.
- [328] L. Torresani, V. Kolmogorov, and C. Rother, “Feature correspondence via graph matching: Models and global optimization,” in *European Conference on Computer Vision (ECCV)*, 2008, pp. 596–609.
- [329] R. Zass and A. Shashua, “Probabilistic graph and hypergraph matching,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [330] O. Duchenne, F. Bach, I. Kweon, and J. Ponce, “A tensor-based algorithm for high-order graph matching,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.

- [331] A. Sotiras, “Discrete image registration: a hybrid paradigm,” Ph.D. dissertation, Ecole Centrale de Paris, 2011.
- [332] P. Fua, “Model-based optimization: Accurate and consistent site modeling.” *International Archives of Photogrammetry and Remote Sensing*, vol. 31, pp. 222–233, 1996.
- [333] T. Binford, “Visual perception by computer. in: Proc. ieee syst. control.” in *IEEE Conference on Systems and Controls*, December 1971.
- [334] C. Kirbas and F. K. Quek, “A review of vessel extraction techniques and algorithms,” *ACM Computing Surveys (CSUR)*, vol. 36, no. 2, pp. 81–121, June 2004.
- [335] D. Rueckert, C. Sonoda, L. I. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, “Nonrigid registration using free-form deformations: application to breast mr images,” *IEEE Transactions on Medical Imaging*, vol. 18, no. 8, pp. 712–721, August 1999.
- [336] T. W. Sederberg and S. R. Parry, “Free-form deformation of solid geometric models,” *ACM Siggraph Computer Graphics*, vol. 20, no. 4, pp. 151–160, 1986.
- [337] P. Fallavollita and F. Cheriet, “Robust coronary artery tracking from fluoroscopic image sequences,” *Image Analysis and Recognition*, vol. 4633/2007, pp. 889–898, 2007.
- [338] W. Neuenschwander, P. Fua, and K. O. Székely, G, *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, Basel, Switzerland, 1995, ch. From Ziplock Snakes to Velcro™ Surfaces, pp. 105–114.
- [339] D. Geiger, A. Gupta, L. A. Costa, and J. Vlontzos, “Dynamic programming for detection, tracking, and matching deformable contours,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, 1995.
- [340] D. Kim, “B-spline representation of active contours,” in *International Symposium on Signal Processing and its Applications*, 1999.
- [341] S. A. M. Baert, W. J. Niessen, E. H. W. Meijering, A. F. Frangi, and M. A. Viergever, “Guide wire tracking during endovascular interventions.” in *miccai*, 2000, pp. 727–734.
- [342] R. von Mises and H. Pollaczek-Geiringer, “Praktische verfahren der gleichungsaufsung,” *ZAMM - Zeitschrift fr Angewandte Mathematik und Mechanik*, vol. 9, pp. 152–164, 1929.
- [343] L. De Lathauwer, P. Comon, B. De Moor, and J. Vandewalle, “Higher-order power method - application in independent component analysis, proc. intl

- symposium on nonlinear theory and its applications (nolta95),” in *Proceedings of International Symposium on Nonlinear Theory and its Applications (NOLTA95)*, Las Vegas, December 1995, pp. 10–14.
- [344] P. A. Regalia and E. Kofidis, “The higher-order power method revisited: convergence proofs and effective initialization,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2000.
 - [345] N. Honnorat, R. Vaillant, and N. Paragios, “Graph-based geometric-iconic guide-wire tracking,” in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2011.
 - [346] A. Sotiras, Y. Ou, B. Glocker, C. Davatzikos, and N. Paragios, “Simultaneous geometric - iconic registration,” in *Medical Image Computing and Computer Assisted Intervention (MICCAI), Part II*, ser. LNCS, T. Jiang, N. Navab, J. P. Pluim, and M. A. Viergever, Eds., vol. 6362. Springer, Heidelberg, 2010, pp. 676–683.
 - [347] U. Kurkure, Y. Le, N. Paragios, J. Carson, T. Ju, and I. Kakadiaris, “Markov random field-based fitting of a subdivision-based geometric atlas,” in *IEEE International Conference on Computer Vision (ICCV)*, 2011.
 - [348] Y. Ou, A. Besbes, M. Bilello, M. Mansour, C. Davatzikos, and N. Paragios, “Detecting mutually-salient landmark pairs with mrf regularization,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2010.
 - [349] U. Kurkure, Y. Le, N. Paragios, J. Carson, T. Ju, and I. Kakadiaris, “Landmark/image-based deformable registration of gene expression data,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 1089–1096.
 - [350] D. H. Ballard, “Generalizing the hough transform to detect arbitrary shapes,” *Pattern Recognition*, vol. 13, no. 2, pp. 111–122, 1981.
 - [351] X. Ren, “Learning and matching line aspects for articulated objects,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
 - [352] K. Siddiqi, a. Shokoufandeh, S. J. Dickinson, and S. W. Zucker, “Shock graphs and shape matching,” *International Journal of Computer Vision*, vol. 35, no. 1, pp. 13–32, 1999.
 - [353] D. Conte, P. Foggia, C. Sansone, and M. Vento, “Thirty years of graph matching in pattern recognition,” *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 18, no. 3, pp. 265–298, 2004.
 - [354] T. Alderliesten, M. K. Konings, and W. J. Niessen, “Modeling friction, intrinsic curvature, and rotation of guide wires for simulation of minimally

- invasive vascular interventions.” *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 1, pp. 29–38, 2007.
- [355] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness (p. 206)*, W. Freeman, Ed., 1979.
- [356] S. C. Narula and C. A. Ho, “Degree-constrained minimum spanning tree,” *Computers & Operations Research*, vol. 7, no. 4, pp. 239–249, 1980.