

CONTENTS

	Page
INTRODUCTION.....	1
0.1 Clinical context	6
0.2 Problem statement	13
0.3 Outline of the thesis	19
CHAPTER 1 LITERATURE REVIEW	21
1.1 Medical imaging segmentation.....	21
1.1.1 The level-set method	22
1.1.2 Level-set segmentation of vascular structures	26
1.1.3 Brain tissue segmentation.....	26
1.2 Vessel detection filters	29
1.2.1 Analysis of the Hessian	30
1.2.2 Optimally oriented flux	31
1.2.3 Polar intensity profile	32
1.3 2D/3D registration of medical imaging.....	33
1.3.1 Types of 2D/3D registration methods	34
1.3.2 Image-to-image methods	35
1.3.3 Model-to-image methods	36
1.3.4 Model-to-model methods	37
CHAPTER 2 OBJECTIVES AND GENERAL METHODOLOGY	41
2.1 Objectives of the research	41
2.2 General methodology	43
2.2.1 Large structures segmentation using the level-set method	43
2.2.2 Vessel detection and segmentation	45
2.2.3 2D/3D registration for the guidance of surgical procedures	46
CHAPTER 3 UNSUPERVISED MRI SEGMENTATION OF BRAIN TISSUES USING A LOCAL LINEAR MODEL AND LEVEL SET	49
3.1 Introduction	49
3.1.1 State-of-the-art techniques and their limitations.....	50
3.1.2 The need for a more comprehensive model.....	53
3.1.3 Organization of this paper	54
3.2 A local linear level set approach to MRI segmentation	54
3.2.1 Formulation of the problem	54
3.2.2 Local linear region model	55
3.2.3 Segmentation using the level set method	57
3.2.4 Extension to a 4 phase model.....	61
3.2.5 Outliers rejection	62
3.2.6 Implementation issues	62

3.2.7	3D extension of the model	63
3.2.8	Computational cost	64
3.2.9	Initialization by Fuzzy C-means clustering	64
3.2.10	Overview of our new segmentation model	65
3.3	Experimental validation	66
3.3.1	Synthetic Data	66
3.3.2	Real Data	67
3.4	Discussion and conclusions	69
CHAPTER 4 3D VESSEL DETECTION FILTER VIA STRUCTURE-BALL ANALYSIS		77
4.1	Introduction	77
4.2	Methodology	80
4.2.1	The <i>structure ball</i> : a local structure model	81
4.2.2	Band-limited spherical harmonics representation	82
4.2.3	Contrast-invariant diffusivity index	85
4.2.4	Geometric descriptors	85
4.2.4.1	Ratio descriptors	86
4.2.4.2	Oriented bounding box	86
4.2.4.3	Fractional anisotropy	87
4.2.4.4	Flux	87
4.2.5	Vesselness measure	89
4.2.6	Multiscale integration	92
4.2.7	Parameter selection	92
4.3	Experimental results	94
4.3.1	Experiments on synthetic images	94
4.3.2	Experiments on clinical images	98
4.4	Discussion and conclusion	99
4.5	Acknowledgement	102
CHAPTER 5 NON-RIGID 2D/3D REGISTRATION OF CORONARY ARTERY MODELS WITH LIVE FLUOROSCOPY FOR GUIDANCE OF CARDIAC INTERVENTIONS		103
5.1	Introduction	104
5.2	Background Information	107
5.3	Translational, rigid, and affine alignment	109
5.3.1	Multi frame alignment	113
5.4	Non-rigid registration	114
5.4.1	Image energy	115
5.4.2	Internal energy	117
5.4.3	Energy minimization	117
5.4.4	Parameter selection	118
5.5	Experimental results	120
5.5.1	Simulations	121

5.5.1.1	Dependence on the initial solution	121
5.5.1.2	Robustness to image noise	122
5.5.1.3	Non-rigid deformation	124
5.5.2	Clinical data	125
5.5.3	Global alignment: evaluation of the performance of the optimizers.....	126
5.5.4	Comparison of the global alignment method with non-rigid registration	129
5.5.5	Global alignment in the multi frame scenario	130
5.5.6	Semiautomatic tracking of the right coronary artery.....	134
5.6	Multimedia Material	135
5.7	Discussion and conclusion	135
CHAPTER 6 DISCUSSIONS		137
6.1	A level-set method using local-linear region models for the segmentation of structures with spatially varying intensity	137
6.2	Brain tissues segmentation in MRI	138
6.3	Accurate local structure modeling and vessel detection using <i>structure balls</i>	139
6.4	2D/3D registration of centerline models with X-ray angiography using global transformation models	140
6.5	Non-rigid 2D/3D registration of centerline models with X-ray angiography	141
6.6	Improved surgical guidance using 2D/3D registration	142
6.7	Application for MAPCAs procedure	143
GENERAL CONCLUSION		145
APPENDIX I EULER-LAGRANGE EQUATION FOR LOCAL LINEAR STATISTICS		149
APPENDIX II SEMI-AUTOMATIC SEGMENTATION OF MAJOR AORTO-PULMO- NARY COLLATERAL ARTERIES (MAPCAS) FOR IMAGE GUIDED PROCEDURES		153
REFERENCES		163

LIST OF TABLES

		Page
Table 1.1	Categories of 2D/3D registration methods	35
Table 3.1	Average Tanimoto index for various segmentation methods on the IBSR 20 normal brain scan	68
Table 3.2	Average Dice index for various segmentation methods on the newer IBSR 18 normal brain scan	69
Table 4.1	Values of the proposed shape descriptors computed on the six vessel voxel configurations presented in Fig. 4.2.	91
Table 4.2	Values of the proposed shape descriptors computed on the eight non vessel voxel configurations presented in Fig. 4.4	91
Table 5.1	RMS residual energy and resulting 2D error, in function of the transformation model.	130
Table 5.2	Total computational time for the two global optimizers.	131
Table 5.3	Mean residual 2D projection error, in mm, calculated after rigid, affine and non-rigid registration for five patients.	132

LIST OF FIGURES

	Page
Figure 0.1 Brain tissue segmentation	8
Figure 0.2 Pulmonary vessels, heart, and lungs	9
Figure 0.3 Normal blood flow through the heart	10
Figure 0.4 X-ray angiography of MAPCAs	10
Figure 0.5 CTA scan of a young child with MAPCAs	11
Figure 0.6 The coronary arteries	12
Figure 0.7 Biplane fluoroscopy of a left coronary artery with CTO	13
Figure 0.8 Research problems	14
Figure 1.1 A simple level-set function	22
Figure 1.2 Evolution of the level-set during a segmentation	23
Figure 1.3 MAPCAs segmentation.....	27
Figure 1.4 Geometry of a 2D/3D imaging system.....	34
Figure 3.1 Two consecutive sample slices from the IBSR, subject 5_8	53
Figure 3.2 Comparison of a segmentation with our local linear model and with the piecewise smooth model	61
Figure 3.3 Results obtained on a slice of IBSR subject 4_8	66
Figure 3.4 Sample segmentations obtained with our method on the BrainWeb synthetic scans	71
Figure 3.5 Results obtained on the 1mm BrainWeb data for various noise level	72
Figure 3.6 Results obtained on the 3mm BrainWeb data, INU=40%.	72
Figure 3.7 Sample segmentation for the IBSR dataset	73
Figure 3.8 Performance index for the 20 normal brain of the IBSR	74
Figure 3.9 Sample result for subject 17_3 of the IBSR database.	75

Figure 3.10	Sample segmentation for the newer IBSR dataset.....	75
Figure 4.1	Icosahedron and subdivision pattern	82
Figure 4.2	Illustration of the shape of the SBall and DBall for various configurations..	83
Figure 4.3	Oriented bounding box of a DBall and computation of the $\mathcal{A}(\mathbf{x})$ function on the circle defined by the \mathbf{v}_3 vector.....	88
Figure 4.4	DBall representation for some interfering patterns.....	90
Figure 4.5	Multiscale response for simple cylindrical patterns with various radii	93
Figure 4.6	Multiscale response for a simple 3D vessel pattern at various angles	95
Figure 4.7	Vesselness filter responses for the slice in Fig. 4.6 for 5 different filters.....	96
Figure 4.8	Vesselness filter responses for the median slice of the structural noise test volume	97
Figure 4.9	True Positive Rate, and False Positive Rate for the fish bone image as a function of the classification threshold	98
Figure 4.10	True Positive Rate and False Positive Rate for the structural noise image as a function of the classification threshold.....	98
Figure 4.11	Vessel detection in a cropped clinical TOF-MRA volume	100
Figure 4.12	Vessel detection in a cropped clinical cardiac CTA volume	101
Figure 5.1	Geometry of the imaging system.....	108
Figure 5.2	Sample images for 2D/3D registration.....	109
Figure 5.3	Progression of a non-rigid registration.....	115
Figure 5.4	Point matching and myocardium constraints	116
Figure 5.5	Total displacement of the centerline as a function of the number of iterations	118
Figure 5.6	Illustration of the behavior of parameter λ in the non-rigid registration method	119
Figure 5.7	Mean 2D and 3D error after non-rigid registration.....	120
Figure 5.8	Creation of the DRRs	122

Figure 5.9	Performance of the optimizers with respect to a perturbation of the initial position	123
Figure 5.10	Performance of the optimizers with respect to the standard deviation of the input DRR noise	124
Figure 5.11	Sample non-rigid registration with simulated data	125
Figure 5.12	Residual 3D error with respect to the simulated non-rigid deformation level	126
Figure 5.13	Sample output obtained with the affine transformation model using nine different optimizers on the Patient 3 dataset	128
Figure 5.14	Mean 2D error, per optimizer, and per patient, after alignment with the affine transformation model.....	131
Figure 5.15	Total computational time, in millisecond, for all local optimizers, when aligning the 3D centerline using the translation only, rigid, and affine transformation models, successively	132
Figure 5.16	Input image, rigid alignment, affine alignment, and non-rigid registration .	133
Figure 5.17	Registration over a sequence of frames	134
Figure 5.18	Tracking the RCA.....	135

LIST OF ABBREVIATIONS

2D	Bidimensional
2D+t	Time series of 2D images
3D	Tridimensional
3D+t	Time series of 3D images
CHU	<i>Centre hospitalier universitaire</i>
CS	Coordinate system
CSF	Cerebrospinal fluid
CT	X-ray computed tomography
CTA	X-ray computed tomography angiography
CTO	Chronic total occlusion
DBall	Diffusion-index ball
DIBCO	Document image binarization contest
DRR	Digitally reconstructed radiograph
dtMRI	Diffusion tensor magnetic resonance imaging
ECG	Electrocardiogram
EM	Expectation-maximization
ETS	École de technologie supérieure
FCM	Fuzzy C-means
FMM	Fast marching method
fMRI	Functional magnetic resonance imaging
FPR	False positive rate
GFA	Generalized fractional anisotropy
GHz	Gigahertz
GM	Gray matter
GMM	Gaussian mixture model

GPU	Graphical processing unit
HARDI	High angular resolution diffusion (magnetic resonance) imaging
HU	Hounsfield unit
IBSR	Internet Brain Segmentation Repository
ICDAR	International Conference on Document Analysis and Recognition
IEEE	Institute of Electrical and Electronics Engineers
INU	Intensity nonuniformity
ITK	Insight segmentation and registration toolkit
LCA	Left coronary artery
MAPCA	Major aorto-pulmonary collateral artery
Max	Maximum
MGH	Massachusetts General Hospital
Min	Minimum
MNI	Montreal Neurological Institute
MPM-MAP	Maximizer of the posterior marginal — maximum a posteriori
MPR	Multiplanar reconstruction
MR	Magnetic resonance
MRA	Magnetic resonance angiography
MRF	Markov random field
MRI	Magnetic resonance imaging
NOFlux	Normalized oriented flux
NSFlux	Normalized spherical flux
OBB	Oriented bounding box
OOF	Optimally oriented flux
PCI	Percutaneous coronary intervention
PFA	Planar fractional anisotropy
PVE	Partial volume effect

RCA	Right coronary artery
RF	Radio frequency
RMS	Root mean square
SBall	Structure ball
SFlux	Spherical flux
SH	Spherical harmonic
SPIE	International society for optics and photonics
SVM	Support vector machine
T	Tesla
TOF	Time of flight
TPR	True positive rate
Voxel	Volume element, a 3D analog to the pixel
VTK	Insight visualization toolkit
WM	White matter

INTRODUCTION

In recent decades, medical imaging had a fundamental impact in almost every branch of medicine. With *in vivo* techniques to image patient tissues, practitioners are able to accurately diagnose an even broader range of conditions, study the structure and the functions of human brain, perform sophisticated laparoscopic interventions on a daily basis, and make detailed surgical plans before entering the operating room. The development of non-invasive imaging procedures, such as magnetic resonance imaging and ultrasound, makes it ethically possible to study healthy subjects, enabling scientists to acquire a deeper understanding of the living body.

But this sophistication is not in itself sufficient to surmount all the medical hurdles. As the world's population is getting older, the increase in prevalence of aging-associated diseases represents a major medical challenge. The causes of many of those conditions, such as the Alzheimer's disease and dementia, are still only partially understood. More prospective and retrospective studies are needed if neurologists are to sharpen their expertise and to gain a better understanding of the causes and characteristics of those conditions. New imaging devices allow to observe an always increasing range of biophysical phenomena, but analyzing the acquired images requires time and energy, which puts a practical bound on how much information we can extract from them.

An aged population is also more at risk of developing vascular disorders. In fact, heart diseases are now the leading causes of mortality in Canada, where they account for 32% of all deaths (Fondation des maladies du coeur, 2008), a situation similar to that in the USA (36.3% of deaths) (AHA News, 2008; Rosamond *et al.*, 2008; Kung *et al.*, 2008) and in the rest of the world (30% of deaths) (World Health Organization, 2008). But vascular problems are not only an issue among the elderly. Heart diseases are among the most commonly encountered congenital disorders, and, among them, are the principal cause of death. In many of those cases, the well being of the patient depends on the outcome of a complex surgical procedure. The always improving imaging tools consistently help to make routine ambitious clinical protocols, but many challenges still lay on the way to risk-free intervention.

While modern medical imaging modalities permit to gain fabulous insight about the patients inner structures, the interpretation and study of medical images remains tedious, complex, and error prone. In some circumstances, the targeted tissues might be poorly visible, or adjacent structures might be hard to discriminate. In other cases, accurately delineating the structures of interest manually might be too time consuming to be of practical use. The problem of accurately positioning the image information within the physical world is also of the highest

importance in many clinical situations. Notwithstanding those various difficulties, recognizing and extracting high level information from medical images is key to gaining a better understanding of a patient inner structures, and, ultimately, will foster the emergence of innovative treatments.

The objective of this thesis is to establish a biomedical image analysis framework for advanced visualization and surgical guidance. It consists of two main parts: 1) the creation of 3D patients-specific models by segmenting some structures of interest in 3D medical scans, and 2) the registration of patient-specific 3D models with other medical images. Image segmentation is concerned with the accurate and repeatable delineation of biological structures in medical images. Such techniques are highly valuable since the segmented image regions constitute a computational model of some important patient structures, and thus enable scientists and practitioners to study the morphology of the tissues in a detailed and quantitative way. Similarly to a road map, a 3D model can also serve as a reference during surgical intervention planning and execution. But it is also possible to envision more sophisticated applications for such 3D model. This model can also serve during the operation to enhance the interventional image. If the 3D model is carefully aligned and warped on top of the interventional images, it will complement those images, and reveals critical information that is not directly available from the interventional methodology. We refer to such alignment and warping process as 2D/3D model to feature image registration, and this constitutes the second part of our framework. Since a 2D/3D registration process positions the 3D model with respect to the interventional image, and thus with respect to the physical world, it creates a convincing and intuitive geometrical reference that can significantly reduce the ambiguity inherent to the guidance of percutaneous interventions.

At the practical level, two broad types of applications were targeted for the evaluation of the performance of the proposed methods: 1) brain morphological studies, and 2) vascular interventions. The study of brain morphology is a fundamental part of many researches on Alzheimer's disease and dementia. At the image processing level, it is often required to discern between the three main tissue classes. In the past few years, this segmentation problem formed an almost canonical benchmark for general purpose and specialized image segmentation algorithms, and thus presented an interesting opportunity to test an important constituent of our framework. In what concerns the vascular interventions, two specific clinical situations were considered. In work done in collaboration with Ste-Justine Hospital, we were interested in segmenting two types of complex vascular structures, major aorto-pulmonary arteries and coronary arteries, in order to ease surgical planning. The assumption being that a 3D model of

the structure is easier to visualize and understand than the full scan. The 2D/3D registration and surgical guidance aspect of framework described here was developed in great part in collaboration with Siemens Corporate Research. There, it was question of registering a model of the coronary arteries with interventional fluoroscopy. Clinically speaking, this can help to reduce ambiguities during one of the most delicate type of non-invasive heart surgeries: percutaneous intervention of chronic total occlusion of a coronary artery. In the following paragraphs, we briefly cover the specific problems that were tackled in this research, and discuss their relative importance. As it will be possible to recognize, the research presented here encompass a large variety of methods and a broad range of practical applications. While this rather large field of view might seem unnecessary at first, we feel that it was essential to clearly demonstrate the theoretical benefits of all proposed image processing approaches.

Segmenting large, relatively uniform, structures in 3D medical scans has many practical applications in biomedical analysis, especially when the exact shape of some tissue is of great interest. Perhaps one of the most striking example of such application is the segmentation of brain tissues in magnetic resonance imaging. In morphological neurological studies, an important task is to separate the white matter from the gray matter and cerebrospinal fluid. Having those structures segmented greatly facilitates the study of the brain morphology and helps to characterize many conditions and pathologies. Unfortunately, various limitations of the acquisition device result in images that are corrupted by different types of noise. Imperfections in the radio frequency coil cause low frequency wave-like interfering patterns to appear and results in non-uniform voxel intensity recordings even for uniform structures. This characteristic of the image is very challenging for most automatic segmentation algorithms since the correct segmentation depends on local feature of the images: no global threshold would produce an acceptable segmentation. This does not imply that the problem cannot be solved by an automatic process, only that a capable method would need to consider more involved processes and features. But what type of process and what kind of features would permit to segment such structures that might be relatively uniform locally, but where the intensity could present large variation from one region of the image to the other, even for the same class of tissues? The solution that is proposed in this thesis is to compute local linear models of the image intensity, which are robust to low frequency non-uniformity, and to use those models to guide the segmentation process. This idea, based on a differential geometry formulation of the problem, has been implemented within the level-set active contour framework. This algorithm was able to generate, to the best of our knowledge, the best published results for this category of algorithms on the publicly available IBSR database.

After having defined our method for the segmentation of large structures, we attempted to adapt this framework to another vastly different problem that is clinically relevant: the segmentation of vascular structures. During work done in collaboration with Ste-Justine Hospital and also with Siemens Corporate Research, it soon appeared that the proposed region based active contour method was of limited use in what concerns the segmentation of very fine tubular structures. The reason of this limitation is twofold. On one side, the viscosity term used in most active contour schemes causes a bias against all small or narrow structures. On the other side, the recorded voxel intensities for very fine structures become dominated by the partial volume effect, which makes regional pixel intensity modelling ineffective at those locations. Those two aspects of the fine tubular structure segmentation problem call for a method that is more specific than the proposed active contour scheme. Fine tubular, or curvilinear, structures can be detected by analyzing the local contrast around each voxel position: if a structure is present, one would expect little intensity variation in one direction, and much larger intensity variation in the other directions. As it would be discussed in chapter 1, many approaches of this type were proposed in the literature, but most of them are ineffective when the target structure present bifurcations or crossings. This brings the following question: is it possible to define a generic image filter that would detect fine curvilinear structures in 3D medical images without excluding bifurcations? In our research, we investigated a new computational model that uses finite differences and a spherical sampling scheme to this end. The results is a new curvilinear structure detection method that is both more specific and more sensitive than the existing method. This new technique has been demonstrated on MRI and CT scan for the detection of different curvilinear structures, and is not tied to any specific 3D imaging modality or biological structures.

This new curvilinear detection method can be used standalone to segment simple tubular structure, or it can complement the segmentation generated by a more general purpose algorithm. For example, for the segmentation of a pulmonary artery tree, the main vessels could be segmented by an active contour method, and the fine ramification by the curvilinear detection method proposed. If the structure of interest is approximately tubular and reasonably salient, the two segmentation results can be merged and will integrate gracefully. Thus, the two methods can be perceived as complimentary. We believe that in a clinical situation, those two techniques can be used together to segment a large variety of vascular structures.

Segmenting tissues in medical images permits to clearly distinguish the target structures from the background elements and ease various morphological studies. In fact, the segmentation result can essentially be regarded as a patient-specific 3D model of the anatomical structures

of interest. If we are to consider the domain of image-based surgical assistance, this brings the question of quantitatively relating the virtual 3D model with the interventional images, and thus with the physical world. That is, how to position the 3D model with respect to the patient body and surgical instruments in the operative room?

Let us assume that during a certain surgical procedure, the interventional biplane fluoroscope has been dutifully calibrated, then it is reasonable to admit that the visualized image are an accurate geometric representation of the scene and that they can be used to guide the surgical operation. Unfortunately, although the interventional images might be correct, there are situations where they are too limited to unambiguously guide the procedure. For example, the visibility of some structures of interest might depends on the presence of a toxic contrast agent, which, for patient safety sake, might only be injected parsimoniously into a human body. In other situation, the 2D nature of the interventional image can make it very challenging to correctly understand the 3D geometry of the target structures. In such cases, displaying a 3D model of the patient structures on top of, or along with, the interventional images can certainly help the practitioner building a more accurate mental image of the patient body, in less time. However, the 3D model needs to be carefully aligned within the scene in order to minimize the observed inter-modality discrepancy and to constitute a coherent representation of the whole scene. Using the calibration of the apparatus as a starting point and the 2D and 3D images features, it is possible to convincingly align the 3D model with the 2D interventional images, a process known as a 2D/3D registration. In mathematical optimization terms, the 2D/3D registration problem is generally considered as well posed if only rigid, affine, or other restricted parametric transformation models are considered. Nevertheless, the rotational nature of the problem makes it strongly non-linear, and thus hard to minimize. In addition, various patient movements, such as respiration and the beating of the heart, can significantly deform the structures of interest during the course of the intervention. This makes it very important to non-rigidly deform the 3D patient model to preserve visual consistency. This, however, is a hard ill-posed problem. Since, surgical guidance is considered, there is also the question of computational complexity. Indeed, it is hard to believe that the surgeon would place an intervention on hold to accommodate a slow registration algorithm. Thus, the question is: is it possible to create a 2D/3D non-rigid registration process that is fast enough to be used during an intervention? The algorithm defined in this thesis takes benefit of a differential geometry formulation to tackle this challenge. This allows the proposed approach to compute a non-rigid 2D/3D registration in typically less than 3 seconds on contemporary commodity hardware. The method has been tested on datasets from 5 patients with cardiac disorders with good results. Thus, its capability, performance and speed makes this method suitable for intra-operative usage.

Taken together, the level-set segmentation method, the curvilinear structure detection technique, and the 2D/3D registration algorithm defined in this thesis form a segmentation and registration framework that is useful in the context of general biological structure visualization, with a special emphasis on vascular interventions. Taken individually, they allow to perform their own range of specific task, as described in the following chapters. Nonetheless, before developing the theoretical aspect of this research further, it is essential to consider the clinical context of the work, and to highlight the challenges arising at the practical level. Then, the research problems that are considered can be stated and discussed more formally. Finally, an outline of this thesis is presented.

0.1 Clinical context

The practical value of any computerized biomedical analysis tool will strongly depends on the specific characteristics of the problem considered. It is therefore necessary to precise the application domains of the research, both in clinical and in biomedical image analysis terms. As introduced earlier, during the course of our research, we were interested in brain tissue imaging and also in two severe vascular conditions: residual major aorto-pulmonary collateral arteries (MAPCAs) in pediatric cardiology, and chronic total occlusion (CTO) of the coronary arteries in adults. In the case of brain imaging, the modality of choice is the magnetic resonance imaging (MRI). In what concerns vascular conditions, the pre operative assessment modality of choice is generally the computerized tomography (CT), a technology that uses X-rays to generate a 3D image of the vascular structures. These CT scans are used to analyze the problematic structures at the diagnostic, planning and surgical levels. In the following paragraphs we expose those various pathological conditions and discuss the challenges associated with the imaging modalities involved in both their study and treatments.

The morphology of the human brain changes widely during the lifespan of an individual (Sakai *et al.*, 2011). In fact, neurologists have observed that the accumulation of memories and experiences are linked with physiological changes in the brain tissues. While in the past this knowledge could only be acquired from the study of post-mortem organs, the advance of MRI in recent decades made it possible to acquire in-vivo images of a patient's brain with virtually no risk for the health of either the patient or the practitioners. The acquisition of multiple snapshots of the brain of an individual at different points in time, in a longitudinal study, allows to appreciate how brain development is related with morphological changes with an interesting accuracy. While some of those morphological changes are part of the normal aging process,

others are linked with the progressions of divers neurological conditions such as Alzheimer's disease, dementia with Lewy bodies, or Parkinson disease.

Although these aging-related neurological disorders are widespread among the elderly, their complex pathophysiological mechanisms are still only poorly understood. As a result, few treatments exist for those conditions (Venneri, 2007). Brain volumetric studies permit to support diagnostic, help monitor the progress of the condition, and permit to gain knowledge about the structural aspects of the diseases. (Bozzali *et al.*, 2008) In a brain volumetric study, the volume and shape of the various brain structures are analyzed and compared. This allows to localize abnormalities, and to quantify the change in volume of the various brain tissues on a voxel-to-voxel basis. The validity and accuracy of the study thus depends on the availability of good quality segmentations of a number of brain scans.

Most brain volumetric studies use T1-weighted MRI as the input modality. Such images present a high soft tissues contrast that, in the best case scenario, allows to clearly identify the boundaries between the congregation of neurons, or gray matter (GM), the interconnections, or white matter (WM), and the cerebrospinal fluid (CSF). Brain tissues segmentation is generally performed initially using a manual or semi-automatic method (Bozzali *et al.*, 2008). When a segmentation becomes available for a certain patient, registration based approaches can then be used to capture relatively small changes.¹ Unfortunately, T1-weighted MR images are generally corrupted by low-frequency intensity non-uniformities, caused by imperfection in the radio-frequency coil of the MRI apparatus (van Leemput *et al.*, 1999). This defect, often referred to as the *bias field* of the system, makes it difficult to clearly distinguish between the different brain tissue classes. In addition, the limited resolution available from the imaging technology implies that more than one tissue class might be present in a single voxel. This so-called *partial volume effect* blurs the boundaries between the different tissues and complicates the task of the neurologists further. The presence of the bias field and the partial volume effect makes it difficult to obtain accurate and repeatable brain segmentation, which in turn can limit the accuracy of brain volumetric studies.

Theses difficulties call for the development of automatic brain tissue segmentation methods that would be more repeatable and requires less effort than the manual or semi-manual methods. However, designing a method that is appropriated for this task is especially challenging because of the presence of intensity non-uniformity. In the worst cases, the voxel representing the white

¹Registration-based brain tissue segmentation is a rapidly evolving field with always improving methods able to cope with larger and larger variations. However, the problem of producing good initial segmentations is still of interest in the definition of atlases, for immature brains, and for brains with abnormalities.

matter in a certain image region are darker than the voxel representing gray matter in another image region, a situation that would confuse most classical segmentation algorithms. In such case, the voxel intensities need to be interpreted locally: on a T1-weighted image, white matter voxels are always lighter than surrounding gray matter voxels, in a certain neighbourhood, even if that is not true for all the voxels over the complete image domain. In addition, the bias field induced variation is generally smooth and of low frequency. This suggests that an appropriate segmentation method could use a region model capable of taking into account such characteristics. Also, since the different biological structures are mostly continuous, a certain degree of spacial coherence can be taken into account. In summary, a segmentation method that is robust to progressive change in the intensity of the structures to segment and that can enforce a certain spacial coherence could help to produce repeatable and accurate brain tissue segmentations.




Figure 0.1 Sample brain tissue segmentation. a) expert segmentation, b) intensity non-uniformity in a white matter region, *the numbers indicate the intensity level*, and c) one case where WM voxels are darker than GM voxels.

The problem of segmenting structures with intensity variation is frequently encountered in medical imaging, and from a broader perspective, this thesis focuses on researching computational tools for diagnosis and surgical planning and guidance. As such, it is our goal to propose methods that are applicable to a large set of medical situations. Consequently, we did not restrict ourselves to brain volumetric studies as an application, and we also address the topic of surgical guidance in vascular interventions. Specifically, we were interested in two other clinical problems: the treatment of MAPCAs, and chronic total occlusions (CTO) of the coronary arteries.

MAPCAs, although part of normal embryonic development, are often harmful after birth, and are associated with severe cyanotic congenital heart defects, such as the tetralogy of Fallot or pulmonary atresia with ventricular septal defect. In tetralogy of Fallot, the pulmonary artery (see Figure 0.2) can be severely atrophied, or even be absent, which prevents oxygen-poor blood from being re-oxygenated by passing through the lung, as is normally the case (see Figure 0.3). Instead, minimal circulation is ensured by one or more MAPCAs directly linking the aorta to the lungs. As the normal function of the aorta is to deliver oxygen-rich blood to the circulatory system, the oxygenation process is inefficient. The treatment of such a condition generally requires one or more delicate surgical interventions. During the unifocalization procedure, the MAPCAs are consolidated in a way that reconstruct a functional pulmonary vascular bed. In addition, any potential secondary irrigation paths are shut in order to restore a normal blood flow. The shape and disposition of the MAPCAs are complex and vary widely from one individual to another. Some vessels can also be very small. Therefore, good surgical planning is crucial to the success of the operation.




Figure 0.2 Pulmonary vessels, heart, and lungs. Frontal (*left*) and dorsal (*right*) views. Illustrations taken from Gray (1918).

Although analysis of blood flow is usually performed using 2D fluoroscopy with targeted contrast agent injection using catheters (see Figure 0.4), understanding the complex geometry of MAPCAs requires a 3D CT angiography (CTA) in the vast majority of cases. Generally CTA, and not magnetic resonance angiography, is favored in these cases, because its higher spatial resolution makes it possible to capture smaller vessels. In addition, because of the CTA's much shorter acquisition time, it might not be necessary to sedate the patient, and so avoid the risk of harm to fragile pediatric patients.




Figure 0.3 Normal blood flow through the heart. The aorta is highlighted in red and the pulmonary artery in blue. Illustration adapted from Yaddah (2006).




Figure 0.4 X-ray angiography of MAPCAs. Contrast agent injection in the aorta clearly demonstrate the passage of the blood through collateral arteries. Lateral (*left*) and frontal (*right*) views.

The 3D information gathered by the CTA scan makes it possible to resolve most geometric ambiguities with good precision and helps to determine the dimensions of the structures, enabling the physician to prescribe the most appropriate treatment. However, since the 3D volume is dense, the structures of interest cannot be visualized directly. Hounsfield units (HU) to intensity and HU to opacity mappings are generally used to generate a selective visualization that

better highlights the structures. Still, it may not be possible to distinguish the structures that are of interest from other structures with similar tissue characteristics using this basic technique. In addition, in some situations an incomplete representation or a poor visualization might result from the uneven propagation of the contrast agent, partial volume effect, or motion artifacts, for example. A high quality patient specific model can be created from the 3D scan which captures the geometry of all the structures of interest and removes any distracting background (see Figure 0.5). Such a model can facilitate the practitioner’s understanding of the patient geometry, which is key to the success of the procedure.




Figure 0.5 CTA scan of a young child with MAPCAs. Volume rendering (*left*), volume rendering with transparent soft tissues (*middle*), and 3D presentation of the segmented aorta and attached MAPCAs (*right*).

MAPCA procedures are usually highly invasive. Surgical guidance is achieved using mono- or biplane fluoroscopy with periodic contrast agent injection. The fluoroscopic images are acquired on-demand, rather than continuously, to minimize patient exposure to X-rays. These 2D images are difficult to interpret and can be geometrically ambiguous. As a result, the surgeon regularly needs to refer to the 3D model during the intervention, but establishing correspondence between the 3D model and the 2D images can be challenging. Keeping the 3D model aligned with the interventional image would help the surgical navigation process, thereby decreasing the risk and possibly reducing patient exposure to X-rays and contrast agent. However, the creation of an accurate patient specific 3D model is difficult and time consuming, and aligning it with the interventional images is a challenging correspondence problem. These issues were important factors in our decision to undertake this research.

Other vascular conditions bring different but related challenges. Specifically, we also address the problem of surgical guidance during the treatment of chronic total occlusions (CTO) of the




Figure 0.6 The coronary arteries. As presented by Gray (1918) (*left and middle*), and a top view of the volume rendering of a CTA scan.

coronary arteries. CTO is the result of a severe accumulation of plaque in the coronary arteries (see Figure 0.6), and are characterized by very low perfusion level over a period longer than three months (Gade and Wong, 2006). This type of lesion is encountered in 15% to 30% of the patients referred for coronary angiography and the presence of CTO is the most powerful predictor of referral for coronary bypass surgery (Grantham *et al.*, 2009). CTO has been referred to as the final frontier in interventional cardiology (Stone *et al.*, 2005). The treatment of CTO generally involves minimally invasive (laparoscopic) percutaneous coronary intervention (PCI). Compared to MAPCAs interventions, although the clinical context is very different, there are many similarities to be found at the image processing level. X-ray fluoroscopy with direct contrast injection is the modality of choice for the guidance of percutaneous coronary interventions. This modality has two major drawbacks: the limited depth perception inherent in 2D images, and the fact that the contrast agent is quickly washed away by the circulating blood. Furthermore, some parts of the coronary arteries are almost totally blocked in CTO. As a results, the contrast agent cannot reach the distal segments of the vessel, making them almost invisible on interventional imagery, as illustrated in Figure 0.7. This situation makes the surgical procedure particularly hazardous, because the practitioner has to guess the position and shape of the arteries, and move the catheter along with great care, so as not to perforate the vessel.

In contrast to X-ray fluoroscopy with direct contrast agent injection, CTA with intravenous contrast injection has different imaging characteristics, and it is often possible to segment the coronary arteries on the 3D scan. Providing precise geometric information can help the physician to make a more accurate diagnosis and prepare a better surgical plan. The extracted model can also serve as a reference during the intervention. However, aligning the artery model with

the interventional images is difficult to achieve, either mentally or with a computational tool. In addition, the coronary artery geometry is altered by both the respiration and heart beat of the patient. Since the CTA is usually acquired under a breath hold and the intervention is performed under free breathing conditions, non-rigid registration is essential to obtaining a good fit between the model and the interventional images. At the image processing level, some of the challenges in the PCI of CTO are similar to those that arise in the treatment of MAPCAs. In both cases, it is crucial to gain an understanding of the geometry of the complex structures of interest. Also relating this geometry to the interventional images is key to the success of the procedure, but performing such a quasi-intermodal registration is a difficult correspondence problem.




Figure 0.7 Biplane fluoroscopy of a left coronary artery with CTO

0.2 Problem statement

Multiple problems arise during visualization of complex anatomical structures in the context of diagnostic and surgical guidance. From an engineering point of view, it is possible to discern problems at three different levels: 1) clinical, 2) medical imaging, and 3) image processing, as summarized in Figure 0.8.

At the clinical level, the practitioner needs to gather the information that is necessary to accomplish his task. So, it is question of observing and measuring the physical world. Problems arising at this level depend only on the physical condition of the patient and on the clinical protocol, and are not tied to a specific imaging modality. In this research, we considered vastly

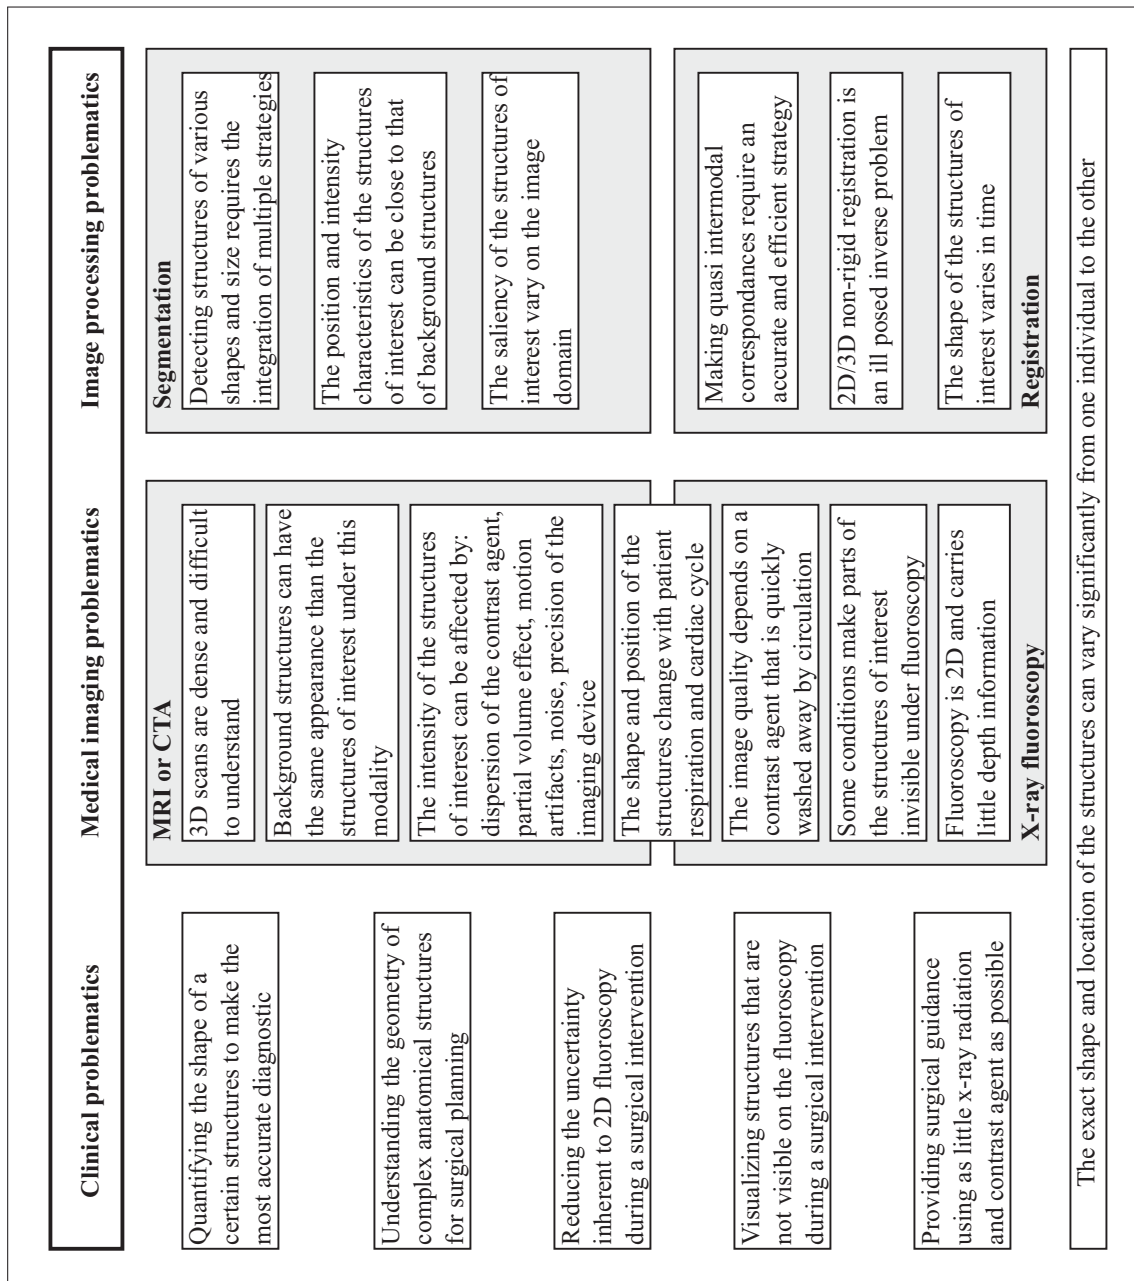


Figure 0.8 Summary of the research problems discussed in section 0.2

different clinical situations. Although they first seem unrelated, we will see later that they share a common set of problems at the image processing level. Specifically, in brain morphological studies, two important questions arise at the clinical level: a) How can the shape of internal anatomical structures be quantified? b) How can the geometry of complex structures be visualized and understood. One way to answer those questions is to delineate the different brain

tissue classes in a 3D scan. Once this is done, a 3D computational can be created, visualized and studied quantitatively.

It is obvious that the two questions presented above are equally important in the context of surgical planning and guidance since dimensioning a biological structure and understanding its shape is important in order to prepare a good surgical plan. However, considering the fact that vascular interventions are generally guided by X-rays raises other crucial questions: a) How can the geometrical uncertainty inherent in limited interventional imagery (e.g. 2D fluoroscopy) be reduced during the guidance of surgical procedures? b) Are all structures of interest visible on the interventional images? If not, how can those structures be visualized during surgical procedures? c) Is it possible to provide surgical guidance with less X-ray irradiation and contrast agent injection, hence optimizing the benefit–risk ratio? In a difficult case, it is critical that the practitioner has access to the best possible information to maximize the chance of a successful procedure. In summary, the practitioner must be able to conceive an accurate mental image of the target structures, and, during the procedure, appropriate navigation guidance must be used to provide live update about the position and shape of the structures.

Various radiological techniques are available that readily offer partial answers to many of these questions. Ideally, the imaging modalities would have infinite temporal and spatial resolution, very high contrast, and would provide the same level of information to the practitioner as direct inspection. In practice, however, because of equipment quality and theoretical limitations, the situation is more complex. This means that problems arising at the medical imaging level interfere with the ability of radiologists to extract the information needed by the practitioner from the images. MRI is the modality of choice for brain imaging since it has better soft tissue contrast than the alternatives. However, the limited precision of the radio frequency coils results in the presence of a significant bias field affecting the image intensity and the partial volume effect blurs the boundary between tissue classes. In addition, a whole head scan is generally performed, which makes discriminating the different brain tissues tedious. In what concerns the two vascular conditions discussed above, CTO and MAPCAs interventions, CTA is used to capture the 3D geometry of the structure of interest before the procedure. This modality makes it possible to accurately capture the patient anatomy in 3D, with some limitations. First of all, since the 3D scans are dense, individualizing a specific structure is difficult. In addition, background structures can have the same appearance as the structures of interest, and the intensity of the structures of interest can be affected by many factors, for example a) the dispersion of the contrast agent, b) the partial volume effect, c) motion artifacts, d) the presence of noise, and e) the precision of the imaging device. Finally, a more fundamental problem is that the

shape and position of the structures change with patient respiration and cardiac cycle. Because of this, the images of the structure observed in the operating room will ineluctably be different than what can be seen on the CT.

In treatment of MAPCAs and CTO, surgical guidance is commonly ensured using X-ray fluoroscopy with direct contrast agent injection, a modality that also has important limitations. In practice, for vascular interventions, the visibility of the structures of interest depends on a contrast agent that is quickly washed away by circulating blood. Moreover, some physical conditions of the patient can make parts of the structures of interest invisible under fluoroscopy, as in CTO. This situation can be precarious because the surgeons have to guess the position and shape of the structure during the intervention. The fact that X-ray fluoroscopy is inherently 2D and has very little depth information is significant: this further complicates the interpretation of the geometry during the procedure. This problem can be alleviated by comparing the interventional images with the pre-operative CTA, but since the shape and position of the structures change with patient respiration and cardiac cycle, this is not a trivial task. Without doubt, these problems can be addressed in different ways. In the future, it is possible that more sophisticated non-invasive operational modalities, such as ultrasound or MRI devices, but with higher spatial resolution, faster acquisition time, and better tissue discrimination properties, would appear and improve the guidance process with very low risk for the patients. However, these new modalities might not be available in all clinical settings, where the current duo of CTA and X-ray fluoroscopy is now well established. In addition, the fact that X-ray fluoroscopy is the most widespread interventional modality ensures that it will remain relevant for many years. In this respect, solutions at the image processing level are more likely to have a short-term impact, and fundamental advances might well be appropriate with the new modalities, when they become available.

The problems related to the visualization of the structures in the 3D scans can be addressed by an appropriate segmentation method, assuming reasonable calibration of the intrinsic parameters of the apparatus. However, the definition of a segmentation method is made more complex by limitations of the medical imaging modalities and by fundamental image processing problems associated with the chosen imaging modalities. For example, extracting the structure of interest from any 3D scan is difficult, because background structures can have the same intensity. Also, various factors can alter the spatial distribution of the intensity of the structure in the volume: dispersion of the contrast agent, partial volume effect, motion artifact, noise, and the intrinsic precision of the imaging device. In what concerns MR brain imaging, the presence of a bias field constitutes the main challenge. This problem is probably not to be

resolved in hardware anytime soon since some of newer MR technologies, such as device with array detectors or the use stronger static field, actually result in a stronger bias field. For vascular structures, especially MAPCAs, the exact shape and position of the structures can vary significantly from one patient to another. At a higher abstraction level, image segmentation is a challenging problem in any practical application. At the same time, this process is essential since it is equivalent to extracting the relevant information from a huge set of data points in order to get an understanding of the physical world depicted by the image. In this problem, it is needed to delineate structures that have a nonuniform appearance over the image domain, variable sizes and various saliency characteristics. These issues are common to most 3D medical imaging applications. For instance, in MAPCAs the intensity of the aorta varies with the dispersion of the contrast agent, while in brain MRI the intensity of the volume varies because of limitations in the imaging device, but in both cases the challenge is to segment complex structures with a spatially variant intensity. In addition, many medical imaging segmentation problems are strongly multiscalar in nature. The fact that the structures under observation can be close to other unrelated structures adds to the challenge as well. For example, with MAPCAs, the aorta is many times larger than the smaller arteries. It is also close to other salient structures, such as the chambers of the heart, which are also filled with contrast agent.

At this point it is important to recognize that the potential impact of 3D medical image segmentation goes beyond the context of discussed clinical situations. The problem of observing and quantifying the shape of anatomical structures is crucial in many areas of medicine. For example, 3D segmentation methods can be used to capture patient geometry for the design of prosthesis, also segmented guts or airways allow to generate virtual endoscopy, thereby easing the diagnostic of cancer, and 3D models of the brain tissues are quantified and analyzed in cross-sectional and longitudinal studies on human brain aging. In vascular interventions, a patient specific 3D geometric model could certainly serve as a useful reference before and during the procedure.

In what concerns surgical interventions guidance, this task is made complex by many of the problems revealed earlier. The information provided during the intervention by the ubiquitous X-ray fluoroscopy is incomplete, due to its 2D nature, and its quality varies over time. For some complex pathologies, like CTO, portions of the structures of interest may be invisible under X-ray fluoroscopy, because they cannot be reached by the flow of the contrast agent. Given adequate knowledge of the geometry of the structure, perhaps obtained using the segmentation method discussed in the above paragraph, those problems could be addressed by registering a 3D model of the structure with the interventional images. In turn, the computed registration

transformation could be used to overlay the 3D model on top of the 2D images, producing enhanced imagery. This combined imagery can facilitate the guidance of surgical intervention since this enable the practitioner to see a reference geometric model in the same frames as the interventional images, which reduce ambiguities in the interpretation of the fluoroscopy. Although both CTA and fluoroscopy use X-rays as the imaging source, these two modalities differ greatly in nature: the former is a reconstructed image, while the latter is a projective image. This quasi-intermodality makes finding correspondences between the images a difficult task. In addition, patient respiration and heartbeat can cause significant shape difference between the 3D model and the 2D images. Using a non-rigid registration method makes it possible to compensate for the discrepancy, but finding its transformation parameters is a challenging inverse problem.

This thesis focuses on the definition of a visualization and surgical guidance framework that is composed of two important parts: 1) the creation of a 3D patient specific model by segmenting a 3D scan, and 2) the registration of this model with 2D interventional fluoroscopies. As discussed above, the definition of such a framework comes with its own set of image processing problems, which represent the main research issues considered in this thesis. In what concern the definition of a segmentation method appropriated for complex biological structures with spatially variant intensity levels, the main hurdles are as follows:

- i. No single segmentation algorithm is suited to all situations, as the structures of interest are characterized by a wide variability in shape and size, and integrating multiple strategies is required;
- ii. The structures of interest can be close, both spatially and in intensity, to other background structures;
- iii. The saliency of the structures of interest vary on the image domain.

For the registration of the 3D model with the interventional images, the principal challenges are as follows:

- iv. making quasi intermodal correspondence requires an accurate and efficient strategy;
- v. 2D/3D non-rigid registration is an ill posed inverse problem;
- vi. The shape and position of the structures vary in time.

Specifically, we proposed to tackle these problems on three fronts. First, generic image segmentation techniques are investigated in order to identify the best methods for segmenting vascular structures with spatially varying intensity. Second, since the generic segmentation methods generally fail at segmenting the smaller vascular structures, localized techniques specialized for the task of detecting and segmenting small blood vessels are considered. Last, but not least, 2D/3D registration methods applicable in context of surgical guidance are researched.

By addressing these image processing problems, this research is aimed at defining methods that are applicable in a clinical setting, and that provide solutions to interventional challenges.

0.3 Outline of the thesis

The next chapter presents a review of state-of-the-art methods that are relevant to the definition of the proposed framework. After, based on this literature review, the objectives of the presented research are defined, and the proposed general methodology is exposed in chapter 2. The three following chapters present the manuscripts written in response to specific research problematic. The manuscript defining our new method for the segmentation of anatomical structures with spatially varying intensity is presented in chapter 3. A curvilinear structure detection filter for vessel detection in 3D scans is described in chapter 4. The method proposed for 2D/3D non-rigid registration of 3D models with live fluoroscopy is the subject of chapter 5. A global discussion is presented in chapter 6, and finally, we summarize the work accomplished in the thesis in a general conclusion.

CHAPTER 1

LITERATURE REVIEW

This chapter presents a review of state-of-the-art methods related to the proposed 3D medical image segmentation and 2D/3D registration framework for the guidance of surgical procedures. This chapter is divided into three sections that are in line with the three biomedical imaging problems exposed in the introduction. The first section starts with a focus on generic medical image segmentation methods, then covers methods specific to the delineation of vascular structures and finish brain tissues segmentation methods. The second section covers methods specific to the detection and segmentation of fine vessel-like structures. Finally, the last section reviews 2D/3D registration techniques that are related to the context of surgical guidance.

1.1 Medical imaging segmentation

A rich literature on image segmentation has emerged in the past 25 years, framed around a variety of mathematical paradigms. Among the methods that have had the most notable influence on the field, we note the snakes method (Kass *et al.*, 1988), based on differential geometry and a variational framework; the level-set and fast marching methods (Osher and Sethian, 1988; Sethian, 1996), which take advantage of implicit representations of continuous functions; the graph cut methods (Greig *et al.*, 1989; Roy and Cox, 1998; Boykov *et al.*, 2001; Kolmogorov, 2004), involving a type of combinatorial optimization; the random walkers (Grady, 2006) and power watershed methods (Coupré *et al.*, 2011), rooted in graph and probability theory; and methods based on probabilistic and non parametric clustering, such as the mean shift algorithm (Comaniciu and Meer, 2002), the MPM-MAP algorithm (Marroquín *et al.*, 2002), and a regression of a Gaussian Mixture Model (Greenspan *et al.*, 2006). The more fundamental work of Mumford and Shah (1989) contributed to our understanding of the segmentation problem.

In this research, the level-set method was selected as the basis for the proposed segmentation algorithm for several reasons: 1) its generality, 2) its suitability for 2D and 3D implementation, 3) its elegant formulation, and 4) its reasonable speed. In addition, level-set methods have previously been proposed for vascular structure segmentation in medical imaging (Lorigo *et al.*, 2001; Vasilevskiy and Siddiqi, 2002; van Bemmél *et al.*, 2003; Manniesing *et al.*, 2006a). The rest of this section gives background information regarding the most common level-set models, then presents specialized methods for the segmentation of vascular structure, and finally, describes other type of methods that have been proposed in the context of brain tissues segmentation. (Manniesing *et al.*, 2006a)

1.1.1 The level-set method

Osher and Sethian (1988) introduced the level-set method to implicitly solve variational problems occurring in the context of the modeling of fluid dynamics. With their implicit representation, an interface is represented using a surface of higher order, that is, a 3D surface is used to represent a 2D curve and a 4D hypersurface is used to represent a 3D surface — see Fig. 1.1.

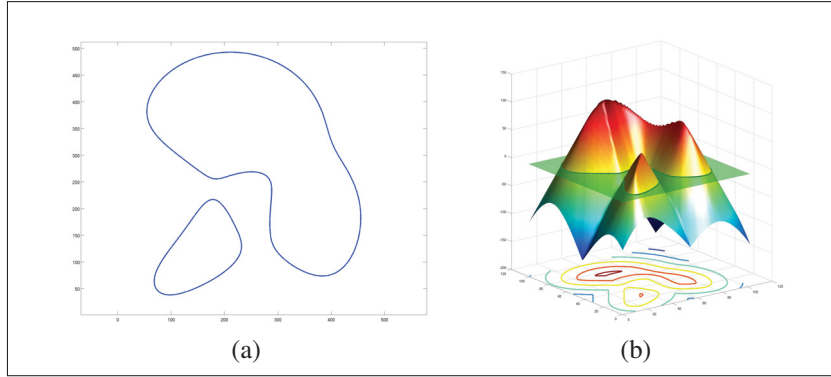


Figure 1.1 A simple level-set function. (a) a set of 2D curves, and (b) its level-set representation.

Formally, a set of curves $C = \partial\omega$ partitioning a domain Ω into two subdomains ω and $\Omega \setminus \omega$ is defined by the 0 level of an implicit level-set function ϕ , defined as follows:

$$\phi(\mathbf{x}) = \begin{cases} +D(\mathbf{x}, \partial\omega) > 0, & \mathbf{x} \in \omega \\ 0, & \mathbf{x} \in \partial\omega \\ -D(\mathbf{x}, \partial\omega) < 0, & \mathbf{x} \in \Omega \setminus \omega, \end{cases}$$

where $D(\mathbf{x}, \partial\omega)$ is the Euclidean distance from a point $\mathbf{x} \in \Omega$ and the curve boundary $\partial\omega$. It is commonly considered that the level-set function ϕ can evolve with respect to an artificial time variable t , and so we note that $\phi(\mathbf{x}) = \phi(\mathbf{x}, t)$ and $\phi_0(\mathbf{x}) = \phi(\mathbf{x}, t = 0)$. In the level-set framework, the essential equation describing the evolution of the implicit function ϕ is

$$\frac{\partial \phi}{\partial t} + \mathbf{v} \cdot \nabla \phi = 0,$$

where \mathbf{v} is a problem specific speed function (Osher and Sethian, 1988; Sethian, 1999).

For practical purposes, with any level-set method, an image segmentation is obtained by: 1) initializing the level-set function ϕ to be a arbitrary surface defined on the image domain, 2) com-

puting the Euler-Lagrange conditions of the selected energy functional, and 3) minimizing the derived Euler-Lagrange conditions. A sample minimization process is depicted in Figure 1.2.

Two important elements compose any segmentation method and must be distinguished: 1) the structure model, and 2) the delineation mechanism. It is clear that a framework such as the level-set method provides that delineation mechanism. However, devising a practical segmentation method requires a set of hypotheses regarding the problem to solve in order to define the structure model. In this respect, most level-set based segmentation methods consider two broad classes of features: edges and region statistics.

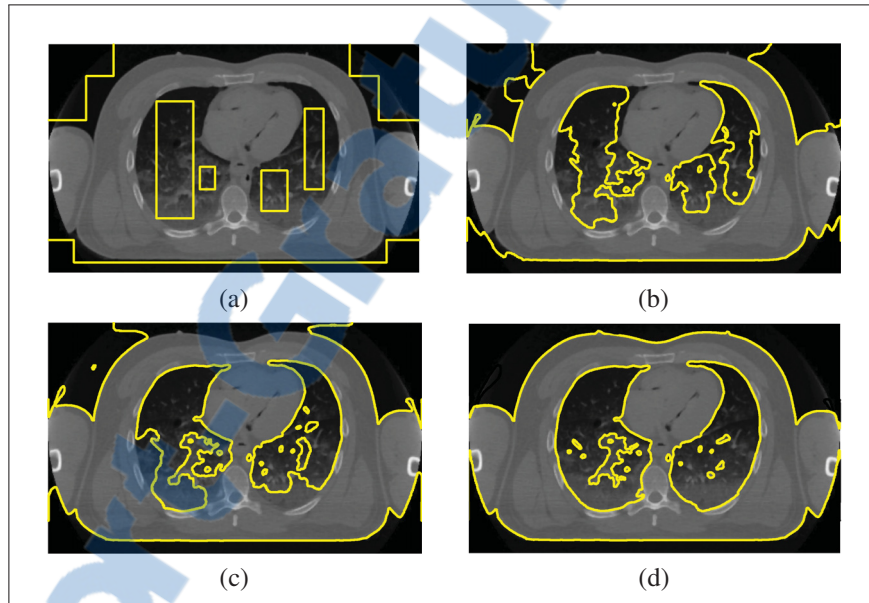


Figure 1.2 Evolution of the level-set during a segmentation. (a) initialization, (b) after 50 iterations; (b) after 100 iterations; (b) final contour, after 300 iterations.

The edge-based geodesic active contour method of Caselles *et al.* (1997) can be considered as an evolution of the snake method (Kass *et al.*, 1988) implemented within the level-set framework. The energy functional they consider is

$$F_{GAC}(u, C) = \int_0^1 |C'(s)|^2 ds + \int_0^1 g(|\nabla u(C(s))|)^2 ds,$$

where $s \in [0, 1]$ is the curve parameter, and the function g is some monotonally decreasing function of the image gradient designed to attract the curve toward the edges, for example $g = 1/(1 + |\nabla(G * u)|^2)$, $G * u$ being the convolution of the input image with some Gaussian filter. The first term seeks to ensure the smoothness of the curve, while the second term is

minimal when the curve lays over an image edge. A limitation of this technique in the context of the segmentation of difficult medical images, such as CT scans, is that weak edges can lead to contour leakage, and in turn to an incorrect segmentation. It is worth noting that Lee and Seo (2006) introduced an optimal formulation of this method, also Malladi and Sethian (1996) previously proposed a related segmentation method.

The active contour without edge of Chan and Vese (2001) makes it possible to solve a restricted formulation of the Mumford and Shah (1989) problem, and so is considered a region-based method. The authors impose the following restriction on the segmentation problem: 1) the number of region is assumed to be known, 2) no steep changes in intensity occur within a region, and 3) the intensity of the region can be considered uniform. Using the previously introduced notation, and letting u^- and u^+ be real numbers representing the mean pixel intensity of the region defined by $\phi < 0$ and $\phi > 0$ respectively, their energy functional for binary image segmentation is

$$F_{ACC}(u, C) = \int_{\phi < 0} (u^- - u_0)^2 dx + \int_{\phi > 0} (u^+ - u_0)^2 dx + \nu |C|,$$

where u represents the input image, ν is a free parameter, and $|C|$ is the length of the partitioning curve. This last term ensure that the boundary is smooth and as short as possible. The first two terms are minimal when the curve delineates two regions of constant intensity. With this method, it is possible to partition an image into approximately constant regions, which limits its use in many practical problems.

Vese and Chan (2002) refined their model by assuming smooth regions instead of constant-value regions. Assuming that u^- and u^+ are now region representatives that vary spatially, the energy function becomes:

$$\begin{aligned} F_{ACS}(u, C) = & \int_{\phi < 0} (u^- - u_0)^2 dx + \int_{\phi > 0} (u^+ - u_0)^2 dx \\ & + \mu \int_{\phi < 0} |\nabla u^-|^2 dx + \mu \int_{\phi > 0} |\nabla u^+|^2 dx + \nu |C|, \end{aligned}$$

where ν is an additional free parameter. The first two terms are minimal when the smoothly varying region model fits the region intensity perfectly. This model is more general than the previous one, but is still oriented toward the segmentation of constant intensity regions. The next two terms favor more uniform region models. A very similar model has been independently proposed in (Tsai *et al.*, 2001). Since the computation of the smooth formulation of

the active contour without edge is relatively costly, various enhanced formulations have been proposed to speed up the computation process (Vese, 2003; Zhang, 2005, 2006).

More specialized level-set region models have emerged in the literature. Gao and Bui (2004, 2005) proposed decoupling the smoothing process from boundary optimization, which allowed for more powerful smoothing. Hongmei and Mingxi (2006) modified the smooth Chan and Vese level-set formulation to use anisotropic diffusion instead of a damped Poisson equation. Statistical interpretation of the Mumford and Shah (1989) functional made it possible to devise methods using more complex statistical regional models such as Gaussian (Brox and Weickert, 2004), linear (Bui *et al.*, 2005), or non-parametric (Kim *et al.*, 2002, 2005). Other authors seek to reunite the edge- and region-based approaches and have proposed hybrid methods (Paragios and Deriche, 2002, 2005; Kimmel, 2003). More recently, (Liu, 2006) introduced a local median region model that is capable of robustly representing smoothly varying regions. A Gaussian version of this method was proposed in (Liu *et al.*, 2007), and a related region-scalable method relying on Gaussian regions is described in detail in (Li *et al.*, 2008b). In this work, the authors compare their method with the piecewise-smooth version of the Chan and Vese level-set method. They demonstrate that, while the two methods compute similar results, theirs is significantly faster.

Researchers have also proposed solutions to the challenges posed by boundary evolution itself. Gomes and Faugeras (2000) proposed a differential equation formulation that preserves the shape of the embedding surface during the evolution of the contours. Li *et al.* (2005a) proposed a different formulation, which also helps to preserve the shape of the level-set function. The Phase Fields method (Rochery *et al.*, 2005; Peng *et al.*, 2008) is an alternative level-set formulation that does not require the assumption of any specific shape for the embedding function. Their formulation is also useful for defining higher-order energy terms driving the evolution of the boundary toward network-like objects, such as roads in a satellite image (Peng *et al.*, 2008). Sundaramoorthi and Yezzi (2007); Sundaramoorthi *et al.* (2009) introduced Sobolev active contours, which allows for the definition of non-shrinking or length increasing contour models. Nain *et al.* (2004) proposed a shape driven active contour flow for the segmentation of elongated structures. Rivest-Hénault *et al.* (2010b) proposed a practical level-set formulation for length increasing contours based on the detection of high curvature regions.

The large variety of level-set methods discussed in the literature gives us the scope to envision many practical applications. However, no method has the best performance in all cases. Devising a usable tool that enables the segmentation of structures composed of heterogeneous regions requires that the technique be adapted to the context of a specific problem.

1.1.2 Level-set segmentation of vascular structures

Some authors have proposed level-set methods specifically designed for the segmentation of vascular structures. The CURVES method of Lorigo *et al.* (2001) is based on the geodesic active contour of (Caselles *et al.*, 1997), but its smoothing term considers only the local minimal curvature, rather than the total curvature. As a result, longitudinal structures are only smoothed along their axes, which prevents them from collapsing on themselves. CURVES is an edge-based method, and so must be initialized close to the desired boundaries.

van Bemmelen *et al.* (2003) combined a formulation similar to CURVE with region-based image terms. Specifically, in addition to an edge based term, they proposed using Gaussian intensity models and the output of Frangi’s vesselness filter as region representatives. For their application, however, they did not find the latter term beneficial. Nevertheless, their level-set model allowed for the accurate segmentation of arteries and veins in MRA scans. Manniesing *et al.* (2006a) proposed a similar model for the segmentation of cerebral vasculature in CT angiography. An interesting aspect of this work is that the authors focus on analysis of the intensity level (in Hounsfield Units) of the structure of interest before defining a Gaussian statistical region model computed on a pre-processed image.

Since these methods use a Gaussian intensity model to represent region intensity, bones, instruments and other salient structures must be masked in a pre processing step in order to compute a practical model of the background. Furthermore, in the context of vasculature segmentation in angiography, the intensity of the structures of interest depends on the presence of a contrast agent, which might be unevenly distributed. Also, the partial volume effect can significantly affect the intensity level of smaller structures. To remedy these difficulties, we propose using a spatially varying region model (Rivest-Hénault *et al.*, 2010a), which is inferred on the basis of user-provided seed points. This model has been applied to the segmentation of the aorta and attached MAPCAs in pediatric cardiology. In comparison to the CURVES model, it allows segmentation of a larger proportion of the structures of interest — see Figure II-5. A related method was proposed shortly after this one by Gao *et al.* (2010) and is available as part of the 3DSlicer medical imaging software¹.

1.1.3 Brain tissue segmentation

Brain tissue MRI segmentation is an almost canonical problem in the field of biomedical image analysis. The main challenge comes from the intensity non-uniformity effect, known as the

¹www.slicer.org




Figure 1.3 Comparison between the segmentation results obtained by the method in (Rivest-Hénault *et al.*, 2010a), images (a)-(c) and (g)-(i), and those obtained using the CURVE equation with a balloon propagation term, images (d)-(f) and (j)-(l). *Top row*: dataset 1, *bottom row*: dataset 2. The gap indicated by an arrow in (h) and (k) is caused by the presence of a catheter.

bias field, that confuse most classical automatic segmentation methods. This bias field usually present itself as a low frequency wave-like interfering pattern and must be taken into account one way or the other if an algorithm is to produce convincing results on brain MRI. Since it was proposed in this thesis to define a segmentation method that is robust to such progressive intensity change, it is essential to covers this important problem. In fact, the local-linear level-set segmentation method presented latter, in chapter 3, was found to be very effective in segmenting brain tissues in MR images, and so brain imaging is the main application in that chapter. Consequently, the state-of-the-art methods that have been proposed for this task are briefly presented in the following paragraphs. Most methods devised for brain tissue segmentation belong to the clustering category. However, a few of the recently proposed methods are based on an active contour framework, and so are more related to the method presented in chapter 3.

Early attempts at automatic brain tissue segmentation involved using generic clustering methods (Worth, 1996) such as MAP (Geman and Geman, 1984), Maximum Likelihood Classification (Duda *et al.*, 2000), and Fuzzy C-means (Hall *et al.*, 1992). The MAP method has been made adaptive and bias-field aware in (Rajapakse and Kruggel, 1998), which resulted in a significant improvement of the results. An Expectation Maximization (EM) method that uses Markov Random Field (MRF) for spatial consistency has been proposed by van Leemput *et al.* (1999), and might be regarded as one of the first truly effective automatic brain tissue segmentation methods. Marroquín *et al.* (2002) proposed a related probabilistic approach, which they claim is more robust to errors in the posterior marginal approximation of the MRF. The related FAST method by Zhang *et al.* (2001) is freely available in the FMRIB Software library. Other clustering methods proposed for the automatic segmentation of brain tissues include: EM classification using various Gaussian Mixture Models (Cuadra *et al.*, 2005; Greenspan *et al.*, 2006; Ferreira, 2007), different bias-field aware Fuzzy C-means methods (Pham and Prince, 1999, 2004; Siyal and Yu, 2005; Bazin and Pham, 2007), classification schemes based on mean-shift clustering (Comaniciu and Meer, 2002; Jimenez-Alaniz *et al.*, 2006; Mayer and Greenspan, 2009), and SVM classification methods (Akselrod-Ballin *et al.*, 2006). In general, methods that include some sort of bias-field modeling outperform the others when used on real-world MRI datasets. Spatial consistency also needs to be considered in the segmentation process, as the expert reference segmentations tend to be relatively homogeneous in a given region.

Level-set active contours handle topological changes naturally and favor spatially consistent solutions, two characteristics that make them attractive for the automatic segmentation of brain tissues. Early proposals (Vese and Chan, 2002; Li *et al.*, 2005b; Angelini *et al.*, 2007) used constant region models that neglected the bias field and were of limited use on real MR images. A user-guided scheme has also been proposed by Gao *et al.* (2010), but it is not suitable in the context of automatic segmentation. A local region model has been proposed by Li *et al.* (2008b) and applied to automatic brain segmentation in (Li *et al.*, 2008a). However, although the authors present impressive results, their method has not been extensively tested on publicly available databases. Also, their method does not exploit the fact that the intra-scan intensity non-uniformity might be different for each class of tissues (van Leemput *et al.*, 1999). In summary, the level-set formulation is attractive in the context of automatic brain tissues segmentation, but current methods need to be formally evaluated. In addition, more accurate region models might lead to better performances.

1.2 Vessel detection filters

Vessel detection and segmentation is a challenging task for a variety of reasons: the structures are small and complex, their observed shape can vary, pathologies may be present that affect their geometry, and contrast is limited in medical imaging, to name a few. At the same time, automatic vessel detection has many potential clinical applications. For example, it can help diagnose vascular disease by highlighting regions with arteriogenesis or stenosis, and facilitate surgical mapping and guidance. Because of the complexity of the structures of interest, many dedicated algorithms have been proposed in the literature. The vessel detection filter methods constitute one type of approach that is of interest. These methods use local features to compute a *vesselness* measure, which represent the level to which a pixel belongs to the *vessel* class. The application of a vessel detection filter on a 3D scan highlights vessel-like structures and discards the background, providing enhanced visualization. Vessel detection filters are also a basic building block in many complex medical image processing applications (van Bemmelen *et al.*, 2003; Sundar *et al.*, 2006; Ruijters *et al.*, 2009; Schneider and Sundar, 2010). In this section, we cover the most important vessel detection methods that can be found in the literature. In addition, dedicated vessel segmentation methods that use related local features are briefly described.

Vessel segmentation methods rely on a variety of techniques to optimize a particular global criterion. This category includes level-set methods (Vasilevskiy and Siddiqi, 2002; van Bemmelen *et al.*, 2003; Law and Chung, 2009), graph cuts methods (Schaap *et al.*, 2009b), particle filters (Florin *et al.*, 2005; Wörz *et al.*, 2009), Bayesian methods (Lesage *et al.*, 2009b), and vessel template matching and tracking methods (Noordmans and Smeulders, 1998; Krissian *et al.*, 2006; Wörz and Rohr, 2007; Friman *et al.*, 2010). A common aspect to all these methods is that they incorporate some sort of vessel-specific local indicator in their definition. Flux operators are the basis of (Vasilevskiy and Siddiqi, 2002; Law and Chung, 2009), Frangi's vesselness filter is used in (van Bemmelen *et al.*, 2003), and Schaap *et al.* (2009b) use local intensity modeling. Various local parametric shape models are also used to describe the structures: the circle (Lesage *et al.*, 2009b; Krissian *et al.*, 2006), the ellipse (Florin *et al.*, 2005; Krissian *et al.*, 2006), the cylinder (Wörz and Rohr, 2007; Wörz *et al.*, 2009), and a user selectable cross-section in (Noordmans and Smeulders, 1998).

Vessel segmentation methods can be considered specific, in the sense that they require either some prior knowledge on the part of the user or an initialization procedure, usually in the form of seed points. In contrast, vessel *detection* filters are usually applied to a significant proportion

of an input volume in voxels and result in a vesselness image highlighting the structure of interest.

1.2.1 Analysis of the Hessian

Certainly the most popular vessel detection filters are those based on the eigenanalysis of the Hessian matrix (Lesage *et al.*, 2009a). For bright tubular structures in 3D scans, the main hypothesis is that the image intensity is mostly uniform in one direction and presents sharp maxima in the other two. Let us define the Hessian matrix $\mathbf{H}(\mathbf{x})$ as the symmetric positive-semidefinite matrix of all second order derivatives at a particular voxel location \mathbf{x} . The three eigenvectors of $\mathbf{H}(\mathbf{x})$ give the principal direction of the second order image intensity variation around \mathbf{x} , and the three eigenvalues indicate the modulus of the variation in those directions. The eigenvalues are interesting descriptors of the local structure and were used in (Lorenz *et al.*, 1997; Sato *et al.*, 1998; Frangi *et al.*, 1998) to define various vessel detection criteria.

Lorenz *et al.* (1997) defines a measure of *line-structure-ness* using the three eigenvalues $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$, ordered by absolute magnitude value, as follows:

$$L = \frac{|\lambda_3| + |\lambda_2|}{2|\lambda_1|}. \quad (1.1)$$

There is no upper bound to L and it will tend to infinity as λ_1 approaches 0, as is the case for the center of vessels and other uniform structures.

A more convenient criterion is proposed by Sato *et al.* (1998), who sort the eigenvalues by *signed* value ($\lambda_1 > \lambda_2 > \lambda_3$) and define their vesselness measure as follows:

$$\mathcal{V}_S = \begin{cases} \exp\left(-\frac{\lambda_1^2}{2(\alpha_1\lambda_2)^2}\right) & \lambda_2 \neq 0, \lambda_1 \leq 0 \\ \exp\left(-\frac{\lambda_1^2}{2(\alpha_2\lambda_2)^2}\right) & \lambda_2 \neq 0, \lambda_1 > 0, \\ 0 & \lambda_2 = 0 \end{cases}$$

where α_1 and α_2 are parameters that can be different for positive and negative value of λ_1 . They motivate this choice by arguing that λ_1 can be slightly positive in the vicinity of a vessel stenosis.

Frangi *et al.* (1998) argued that the previously defined criteria are not strict enough to efficiently discard interfering structures. They introduced a vesselness measure that uses two indicators \mathcal{R}_B and \mathcal{R}_A to distinguish blob-like and plate-like structures respectively, plus a *second order*

structureness term \mathcal{S} that accounts for the saliency of the structure. These indicators are defined as follows:

$$\mathcal{R}_{\mathcal{B}} = \frac{|\lambda_1|}{\sqrt{\lambda_2 \lambda_3}}, \quad \mathcal{R}_{\mathcal{A}} = \frac{|\lambda_2|}{|\lambda_3|}, \quad \mathcal{S} = \sqrt{\sum_{j \leq 3} \lambda_j^2}$$

Based on these expressions, they define their vesselness measure as follows:

$$\mathcal{V}_F = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ (1 - \exp\left(-\frac{\mathcal{R}_{\mathcal{A}}^2}{2\alpha^2}\right)) \exp\left(-\frac{\mathcal{R}_{\mathcal{B}}^2}{2\beta^2}\right) (1 - \exp\left(-\frac{\mathcal{S}^2}{2c^2}\right)) & \text{otherwise} \end{cases},$$

where α , β and c are free parameters. Frangi's vesselness measure is more discriminative than Sato's, but at the same time its performance with very small vessels is reduced. In addition, this method is very sensitive to the presence of vessel junction.

In practice, the second order derivatives of the Hessian matrix are computed using a combination of Gaussian filters. This gives us the opportunity to compute vessels appearing at different scales using the framework of Lindeberg (1996), which is exactly what is proposed in (Lorenz *et al.*, 1997; Sato *et al.*, 1998; Frangi *et al.*, 1998).

1.2.2 Optimally oriented flux

Vessel detection filters based on flux operators are conceptually different. In particular, Law and Chung (2008) consider the total flux passing through a circle positioned perpendicularly to the vessel direction to define their vessel detection measure. Their flux operator, oriented in the $\hat{\rho}$ direction and computed for a radius r , is defined as follows:

$$\text{OOF}(\mathbf{x}, r) = \int_{\partial S_r} \left((\mathbf{v}(\mathbf{x} + \mathbf{h}) \cdot \hat{\rho}) \hat{\rho} \right) \cdot \hat{n} \, dA,$$

where \mathbf{v} is the input image gradient, dA is the infinitesimal area on ∂S_r , \hat{n} is the outward unit normal to ∂S_r at position \mathbf{h} , and $\mathbf{h} = r$. A fast method for computing the optimal direction $\hat{\rho}$ and for estimating the integral is presented in (Law and Chung, 2008). The detection of vessels at multiple scales is performed by computing the OOF for many r values. The authors argue that their method is better than the Hessian-based approach for detecting vessels located close together, because scaling the integrating circle does not blur the structure, as it does with the scale-space framework of Lindeberg (1996).

1.2.3 Polar intensity profile

Qian *et al.* (2009) highlighted the fact that most vessel detection filters are plagued by the assumption of a single vessel direction per voxel. As a result, the vesselness signals of those filters incur high loss at vessel bifurcations. The authors address this problem by proposing a new vesselness measure based on polar intensity profile. The intensity profile they consider is as follows: for an image I voxel at position \mathbf{x} , the intensity profile in a certain direction $\theta = (\psi, \phi)$ is the profile of all voxel intensities along a ray released at \mathbf{x} , oriented in θ and with a length r . This 1D intensity profile is computed in many directions. The hypothesis is that the intensity profile has low variability in the direction of a vessel, and high variability in the other directions. They compute a polar intensity profile by dividing the sphere volume into 3D sectors defined by directions $\theta \pm (\Delta\psi, \Delta\phi)$ and radius coordinates $[0, r]$. Practically speaking, the 3D sectors S_θ are composed of many sampling points \mathbf{u} . For each 3D sector, an intensity variation metric is computed:

$$\text{Dev}(\mathbf{x}, \theta) = \frac{1}{|\#S_\theta|} \sum_{\mathbf{u} \in S_\theta} (I(\mathbf{u}) - I(\mathbf{x}))^2,$$

where $|\#S_\theta|$ is the cardinality of the sampling point set for sector S_θ . From there, they define the probability $p_v(\mathbf{x}, \theta)$ of having small intensity variation within the discretized orientation region S_θ as $p_v(\mathbf{x}, \theta) = c \exp(-\beta \text{Dev}(\mathbf{x}, \theta))$, where c and β are positive real valued parameters. An entropy based measure, $v(\mathbf{x})$, is defined to evaluate the *tightness* of the polar intensity profile:

$$v(\mathbf{x}) = \exp(-\tau H(p_v)) = \exp\left(\tau \int p_v(\mathbf{x}, \theta) \log p_v(\mathbf{x}, \theta) d\theta\right).$$

Since this quantity gives little information about the saliency of the structure, they also propose a measure for a locally bright structure, which is:

$$b(\mathbf{x}) = 1 - \exp\left(-\frac{\mu^2(I_{S_{\theta_{\min}}(\mathbf{x})})}{2\gamma}\right).$$

Here, $\mu^2(I_{S_{\theta_{\min}}(\mathbf{x})})$ is the mean intensity within the region with the minimum deviation. Finally, their $[0, 1]$ normalized polar profile vesselness measure is

$$\text{PPV}(\mathbf{x}) = b(\mathbf{x})v(\mathbf{x}).$$

This measure has very good performance at vessel bifurcations, but we have found it sensitive to the presence of interfering structures in our experiments — see chapter 4.

To summarize, vessel detection filters are valuable medical imaging processing techniques, and they serve as a basis for many more advanced methods (van Bemmél *et al.*, 2003; Sundar *et al.*, 2006; Ruijters *et al.*, 2009; Schneider and Sundar, 2010). However, the performance of the most commonly encountered Hessian-based vessel detection filters is limited at vessel junctions because of an implicit single vessel assumption that is built into the method. The oriented flux operator and the polar intensity profile vesselness filter perform better in this respect, but, as pointed out above, our experiments (presented in chapter 4) show that they are sensitive to the presence of interfering structures. In addition, the polar intensity profile of (Qian *et al.*, 2009) cannot easily be saved and reused in a subsequent application — at least not without satisfying a huge memory requirement. Overall, there seems to be a difficult trade-off, between sensitivity and discrimination power with existing vessel detection methods.

1.3 2D/3D registration of medical imaging

Registering medical image modalities consists in placing two datasets in the same reference frame. In 2D/3D registration one modality is generally a 3D pre-intervention CT or MR scan and the 2D modality is almost always some sort of x-ray radiography. For example, in Figure. 1.4 the transformation \mathbf{T} that best aligns the 3D modality reference point (C) with that of the biplane c-arm gantry (X) needs to be computed. In the case of non-rigid registration, not only does there need to be compensation for differences in position and orientation, but the shape discrepancy must also be taken into account.

The registration process generally involves computing the transformation of one dataset into the other by minimizing a particular certain energy function. So, the registration of two medical images involves three main elements: 1) a transformation model, 2) an image comparison strategy, and 3) an optimization algorithm. In the specific context of the 2D/3D registration of medical imaging for intervention guidance, there is an inherent dimensionality problem that must be accounted for by the image comparison strategy. In addition, when considering the registration of CTA or MRA with X-ray fluoroscopy, the process can be considered as quasi-intermodal or intermodal, which means that the intensity level of the two images cannot be compared directly. In this respect, the most critical aspect of any method for the 2D/3D registration of medical imaging for intervention guidance is the image comparison strategy.




Figure 1.4 Geometry of a 2D/3D imaging system

1.3.1 Types of 2D/3D registration methods

It is possible to classify the various 2D/3D registration methods that have been described in the literature essentially by inspecting two of their most important characteristics: 1) what is the basis for comparison, and 2) what is the dimensionality of the comparison domain². The basis for comparison refers to the abstractions that are actually compared during the registration process. We consider two classes of basis for comparison for both the 2D and 3D modalities: image bases and model bases. An image basis involves simply the input images, either the 2D fluoroscopies or the 3D volume, with or without pre processing with some low level filter, such as a gradient or vesselness filter. In contrast, a model basis is composed of some features of interest. For the 2D modality, this basis could be, for example, the thresholding of a vesselness filter response or the result of a structure identification process. For the 3D modality, the basis could consist of a 3D segmentation of the structures of interest, a centerline representation or a more complex model. The dimensionality of the comparison domain refers to the dimension of the similarity measurement space. The comparison domain is either 2D or 3D. Bidimensional comparison domains are associated with projective methods. The key aspect of these methods is that the 3D representation is projected onto the 2D imaging planes as part of the comparison process. The 3D to 2D projection can be simulated X-ray fluoroscopies, called digitally reconstructed radiographs, or a maximum intensity projection. The methodology is reversed when a 3D comparison domain is chosen. With reconstruction methods, it is the 2D image features

²More complex classification scheme are also possible, see (Maintz and Viergever, 1998; Lester and Arridge, 1999; Markelj *et al.*, 2012).

Table 1.1 Categories of 2D/3D registration methods

Basis for comparison		Comparison domain	Representative work
2D modality	3D modality		
Image	Image	2D (projection)	(Birkfellner <i>et al.</i> , 2009)
Image	Image	3D (reconstruction)	(Tomazevic <i>et al.</i> , 2006)
Image	Model	2D (projection)	(Turgeon <i>et al.</i> , 2005; Ruijters <i>et al.</i> , 2009)
Model	Model	2D (projection)	(Benameur <i>et al.</i> , 2003; Groher <i>et al.</i> , 2009)
Model	Model	3D (reconstruction)	(Zheng <i>et al.</i> , 2009)

that are projected into the 3D space by applying a stereographic reconstruction algorithm. The back-projection strategy also leads to 3D comparison, and it consists of measuring the error by directly considering the 3D projection line associated with a 2D point, with no attempt to reconstruct a specific 3D position.

The comparison basis can be different for the 2D and 3D modalities. Also, the dimensionality of the comparison domain is relatively independent of the basis for comparison, although some configurations are more intuitive than others. It is thus possible to classify most proposed 2D/3D registration methods into five classes, as presented in Table 1.1.

The other important aspects of the registration methods mostly derive from the basis for comparison and domain selected. Nevertheless, they may have a critical impact on the overall performance of the method. The comparison metric is a function that makes it possible to measure the distance between modality representatives. This can be a mutual information measure in the case of image-based registration, or the Euclidean distance for model-based methods. The transformation model selected determines the maximal precision that can be achieved by a registration method. Generally, the maximum possible precision increases with the model complexity (Lester and Arridge, 1999). However, added complexity can result in more local minima on the optimization surface, which in turn calls for powerful optimization algorithms. Below, we describe some of the most interesting registration methods that are relevant for the 2D/3D registration of medical imaging for cardiac interventions guidance.

1.3.2 Image-to-image methods

Image-to-image projective 2D/3D registration methods are the most frequently encountered methods in the literature (Markelj *et al.*, 2012). They involve projecting the volumetric information onto the 2D imaging plane. In general, the physical properties of the imaged tissues, as

well as the characteristics of the imaging source, are modeled and encoded into the projection algorithm. In this case, the projected volume results in a synthetic X-ray fluoroscopic image called a Digitally Reconstructed Radiograph (DRR), which is very similar to a real 2D fluoroscopic image (Hipwell *et al.*, 2003). Other projection algorithms, such as maximum intensity projection, have been considered by some authors (Kerrien *et al.*, 1999). In this category of registration methods, the two principal research directions have been the efficient creation of high-quality DRRs (Khamene *et al.*, 2006), and the definition of the best performing comparison metrics (Liao *et al.*, 2006; Birkfellner *et al.*, 2009).

A related registration technique is the image-to-image reconstruction-based 2D/3D registration method of Tomazevic *et al.* (2006), in which a synthetic volume is reconstructed using multiple 2D angiography images acquired with a 3D rotational X-ray device, and the filtered back-projection technique of Grass *et al.* (1999). An asymmetric multifeature mutual information measure is then minimized to compute the registration. A variation of this technique consists of applying a gradient filter on the input image in order to produce a more selective basis for the similarity criterion (Markelj *et al.*, 2008). Approaches of this nature are limited by the necessity of interrupting the intervention procedure to process for rotational acquisition. In addition, they have only been demonstrated on high contrast bone datasets. Nevertheless, development in the field of 3D rotational X-ray techniques are likely to make image-to-image reconstruction-based 2D/3D registration methods more attractive in the future.

A limitation that is common to most image-to-image 2D/3D registration methods is that the contrast in the projection image is lower than the input volume, which limits potential applications to 3D scans with highly contrasted structures: bones in (Khamene *et al.*, 2006; Tomazevic *et al.*, 2006; Birkfellner *et al.*, 2009), and cerebral structures in MRA in (Hipwell *et al.*, 2003). To remedy this difficulty, DRRs can be generated by considering only: *a*) the subset of the 3D volume corresponding to the structures of interest (Turgeon *et al.*, 2005), *b*) a realistic mesh model (Aouadi and Sarry, 2008), or *c*) a reference CT atlas (Hurvitz and Joskowicz, 2008). The creation of such a volumetric model makes it necessary to segment the 3D volume before the registration, but this also makes it possible to produce more highly contrasted DRRs that enable registration of images of soft tissue structures; for example, coronary arteries in (Turgeon *et al.*, 2005).

1.3.3 Model-to-image methods

The problem of the 2D/3D registration of a 3D model of the aorta and attached MAPCAs has been addressed by our team in Julien Couet and *et al.* (2012). In this research, DRRs

were generated using a fast wobbled splatting technique (Birkfellner *et al.*, 2005). A problem that generally occurs with this method is that a flies screen-like pattern might appear on the DRRs because of the magnification induced by the simulated imaging system. We proposed to alleviate this problem by using a triangle convolution filter, which under certain circumstances is equivalent to using a 3D image upsampled by trilinear interpolation (Cosman, 2000). The minimization is driven by an improved hill climbing algorithm. This method resulted in a target registration error of less than 2.060 ± 2.712 mm on a difficult real-world MAPCAs dataset. The main drawback of this method is that the relatively high computational time required makes it more suitable for off-line applications.

Registration methods based on the projection of full volumes or dense models on the 2D imaging plane always involve relatively high computational cost. This places a practical limit on the complexity of the registration transformation models considered. As such, the results reported for many 2D/3D registration methods are limited to rigid transformation (Hipwell *et al.*, 2003; Turgeon *et al.*, 2005; Khamene *et al.*, 2006; Tomazevic *et al.*, 2006; Aouadi and Sarry, 2008; Markelj *et al.*, 2008; Birkfellner *et al.*, 2009) or a parametric approach with a small number of parameters (Hurvitz and Joskowicz, 2008). With a simpler 3D model, the computational cost is greatly reduced. Ruijters *et al.* (2009) introduced a projective model-to-image 2D/3D registration method for the alignment of coronary arteries centerline model with live fluoroscopies. Their method is faster than those in the preceding categories, but still requires several seconds per frame with only a rigid transformation model.

1.3.4 Model-to-model methods

Model-to-model 2D/3D registration methods are potentially much faster during the optimization step. Feldmar *et al.* describe a framework for the registration of points, lines, and surfaces in (Feldmar *et al.*, 1994, 1995, 1997). The registration transformation is first approximated using geometrical analysis and is refined in a second step involving the Iterative Closest Point (ICP) algorithm. They present results obtained on a variety of datasets, including brain vasculature. Their method requires user interaction at the initialization level, which is a limitation for some applications, and they do not clearly describe how the features are extracted from the various images. More recently, Benameur *et al.* (2003) proposed a method of the registration of a deformable 3D mesh model of the vertebra with radiographic images. In this work, the features are automatically extracted in 2D using Canny's edge detector. The registration parameters are computed using a gradient descent approach following an initial alignment step.

An efficient model-to-model 2D/3D registration method for coronary artery centerline model alignment with x-ray images have been proposed by Sundar *et al.* (2006). This method is based on a Distance Transform (DT) of the 2D features on the image plane, a concept previously introduced by Paragios *et al.* (2003) in the context of 2D/2D registration. The main advantages of this method are its relatively large capture range and its computational efficiency. Indeed, by using the DT of the 2D image, the projection error of the 3D model can be evaluated directly, without the need to estimate matched point as in (Feldmar *et al.*, 1995). In this work, the 2D features are extracted using a thresholding of the digitally subtracted angiography. A related method has been proposed by Duong *et al.* (2009), relies on semi-automatic vessel centerline extraction on the 2D images using a vesselness measure and the Fast Marching method (Sethian, 1999).

A graph based non-rigid 2D/3D method for the registration of a vessel centerline in the one-view scenario is presented in (Groher *et al.*, 2009). These authors use a thinning procedure to extract graphs in both the 2D and 3D images, and they propose an original length preserving criterion to ensure the geometrical plausibility of the computed non-rigid registration. However, the high computational time involved makes this method impractical for live applications. An efficient graph-based method has been proposed by Liao *et al.* (2010), who propose using DT and a thin plate spline deformation model (Bookstein, 1989) to greatly reduce the computational cost. However, their method use features specific to the aorta, and so is not suitable for all applications.

The remaining class of 2D/3D registration methods consist of the reconstruction-based model-to-model techniques of Zheng *et al.* (2006, 2009). In these works, the authors focus on constructing a bone surface model from a limited number of 2D X-ray images. Feature points are extracted in the 2D images, and then matched using a non-rigid iterative process. The final reconstruction requires an accurate statistical shape model that is deformed as a function of the established correspondences. Hence, the practicability of the method is limited by the availability of a suitable statistical shape model, and it is still not clear that such a model can be constructed for soft tissue structures that are subject to large variations in shape. Also, at about 2 minutes per reconstruction, the computational cost of this method is relatively high.

Two essential, but somewhat conflicting qualities, of 2D/3D non-rigid registration methods in the context of cardiac intervention guidance are accuracy and speed. The maximal accuracy of a method depends on the transformation model considered. In the context of cardiac intervention, the pre operative CTA is acquired under a breath-hold, whereas the interventional images are acquired in a free breathing condition. This results in significant shape differences between

the two datasets and makes the use of a non-rigid deformation model indispensable. However, using such a model increase the complexity of the problem, and usually results in high computational time, which can make the method impractical. We could not find a method in the literature that completely solves the problem of guidance in challenging cardiac interventions.

CHAPTER 2

OBJECTIVES AND GENERAL METHODOLOGY

2.1 Objectives of the research

As discussed in the introduction, the general objective of this thesis is to define a biomedical image analysis framework for advanced visualization and surgical guidance. This framework consists of two important parts: 1) the creation of 3D patients-specific models by segmenting some structures of interest in 3D medical scans, and 2) the registration of patient-specific 3D models with interventional fluoroscopy for intervention guidance. The literature review presented in the previous chapter highlights some of the limitations of the current medical image segmentation and registration methods. The presented limitations suggest that a common challenge in biomedical imaging is: how to adapt to local structure variation while preserving consistency? In this respect, localized methods based on differential geometry that permit a better adaptivity to the image characteristics have been researched.

Specifically, with level-set segmentation method there is always a trade-off between precision capture range. For example, edge-based active contour models (Caselles *et al.*, 1997; Lorigo *et al.*, 2001) are driven toward the nearest maxima of the gradient modulus image. Therefore, they need to be initialized relatively close to the structure of interest. On the other side, formulations based on constant region model have little dependency (Chan and Vese, 2001)¹ on the initial position of the contour, but the constant region model strongly limit their performance with most real world images. It can be hypothesized that a more flexible local region model could provide a more balanced trade-off between the capture range and the precision of the method. A first specific objective of this thesis is thus:

Specific Objective 1

To define a segmentation model that is robust to low frequency changes in the intensity of the voxels corresponding to the structures of interest.

The literature indicates that global segmentation techniques, such as the region-based level-set methods, are not suitable for the segmentation of very fine vascular structure with a cross-section that can be as small as one or two voxels in diameter. In this context, vessel detection filters can complement another segmentation model (Manniesing *et al.*, 2006b). Unfortunately,

¹Or even no dependency in the case of Rochery *et al.* (2005).

the most popular existing vessel detection filters (Sato *et al.*, 1998; Frangi *et al.*, 1998) suffer from high signal loss in regions close to bifurcations. In contrast, newer methods might not be discriminative enough to be useful on thoracic CTA volumes. Therefore, a second specific objective of this thesis is:

Specific Objective 2

To devise a vessel detection filter for the segmentation of fine vascular structures that is robust to the presence of bifurcations and has good discriminative capacity.

It is believed that the combination of a global segmentation method with local structure descriptors would make it possible to efficiently and accurately segment many vascular structures. The segmented structures can serve at the diagnostic and surgical planning stage, where they would permit a better understanding of the geometry. However, mentally mapping this information to the structure observed on the interventional images is difficult, time-consuming, and error prone. In addition, the uneven quality of the X-ray fluoroscopy, and the structure deformations caused by body functions further complicate this task. To alleviate this problem, a solution is to register a 3D patient-specific model of the structures of interest with the interventional images. Overlaying the model on top of the 2D X-ray images significantly reduce ambiguities in their interpretation. The literature shows that many 2D/3D registration methods have been proposed for similar tasks in related contexts. Nevertheless, no method has been found that is capable of providing a non-rigid registration fast enough to be used during the operation. In this respect, it is not necessary to perform the registration in realtime, a method that run in interactive time would allow to perform many updates during the course of a procedure, after the contrast agent has been injected. Hence, the third specific objective considered is:

Specific Objective 3

To propose an efficient 2D/3D non-rigid registration methodology for the guidance of complex procedures.

The three proposed specific objectives offer a two-step framework for enhanced surgical guidance consisting in 1) segmenting the structure in the pre-operative CTA or MRI scan, and 2) registering this patient specific model with the interventional images. The achievement of the proposed objectives is however not straightforward, and specific considerations are needed at each stage. As such, the next section discusses the proposed methodology and the specific challenges encountered.

2.2 General methodology

The general methodology considered is directly tied to the specific objectives defined above, and thus consists of three principal parts: 1) large structure segmentation using the level-set method, 2) vessel detection and segmentation based on local image structure analysis, and 3) 2D/3D registration for the guidance of surgical procedures. We favored methods based on differential geometry because of their elegant formulation, deep roots in the literature, and long term relationship with computer graphics. At the practical level, they tend to produce smooth outputs that are in accordance with the organic nature of medical images. Each step is briefly outlined here, and is the subject of a complete chapter later.

2.2.1 Large structures segmentation using the level-set method

The segmentation of structures with spatially variant intensity has been identified as a particularly problematic case. It is also an important problem, because such structures can be found in most real world medical images. Hence, a segmentation model is defined that is robust to small changes in the intensity of the voxels corresponding to the structures of interest, a task that is directly linked to the realization of the first specific objective of this thesis. During our research, we investigated the potential of optimization algorithms, acting at a global level, to obtain a segmentation of some anatomical structures. Methods in this category include the geometrically motivated level-set and graph-cut methods, along with their various refinements, and also clustering methods such as the fuzzy C-means and Expectation Maximization algorithms. Our research converged to the level-set approach, which is based on a compelling variational framework, because it has been successfully used for the segmentation of some vascular structures and brain tissues.

As stated in the first objective, our focus was on methods that could efficiently deal with some intensity nonuniformity in the intensity of the voxels representing the structure. At the time we started this research, level-set methods using a localized region model have just been proposed for tackling similar problems. This idea has a great potential: while the contour defined by the level-set still acts at a global scale, the statistical regional model is localized and permits a certain adaptation to the local image contrast. The work presented in chapter 3 pushes this idea further by defining a local-linear model. Not only is fitness of the voxel intensities tested by this new functional, but the fitness of the voxel intensities spatial derivatives is tested as well. This opens the way to model image regions with significant gradients, and is an idea that could certainly be used for the segmentation of vascular structure. However, we found that our new local-linear level-set method demonstrated better potential for use in brain tissue MRI.

Our specific objective, in the context of automatic brain tissue segmentation, is to partition a T1-weighted MRI scan of the head into three different classes: the nerve connections (white matter or WM), the congregations of neurons (gray matter or GM), and the cerebrospinal fluid (CSF), without user intervention. It is generally assumed that the skull and other background structures have previously been removed using another method. Defining an accurate method to automatically segment brain tissues is a difficult task, principally because the intensity of the pixels describing each class is not constant, and can vary significantly from one region of the scan to another. The main sources of intrascan intensity nonuniformity (INU) are limitations of the imaging device itself, such as imperfections in the radio frequency coils. This phenomenon is commonly referred to as the *bias field* of the MRI system.

Over the years, two main avenues of research have been explored to find solutions to the bias field problem: 1) producing unbiased MR images, and 2) developing segmentation methods that are as resilient as possible to the presence of INU. In the first case, both retrospective and prospective methods have been proposed. However, the retrospective methods are limited in the bias field for which they can compensate, and the prospective method are still not been adopted in most current clinical workflows. Consequently, the development of a resilient segmentation method is of interest to practitioners. In the literature, most methods developed to the automatically segment brain tissues are based on clustering methods (EM, MPM-MAP, GMM, fuzzy C-means, Mean Shift, to name a few) designed to explicitly compensate for a bias field that varies slowly. The results produced by the best methods are certainly an improvement over the basic technique, but all methods are limited when applied to lower quality scans with high INU.

In this thesis, we propose a more geometrically oriented method defined in a variational framework. This approach allows us to exploit the spatial consistency of the tissues during the segmentation. Resilience to the presence of a moderate level of INU is implemented implicitly by the use of the local-linear region model briefly introduced above. Our new regional energy formulation is, however, more complex than existing models, and finding a gradient direction that guarantees that the energy is minimized requires a significant amount of computation. For practical purposes, an approximation is used to significantly reduce the computational complexity. This approximation is presented in appendix I, and a validity condition is provided. An experimental validation study has been conducted using both synthetic and clinical MRI data of brain tissues, and detailed results are presented in chapter 3 of this thesis.

The work described here meets the goal of defining a segmentation model that is robust to small changes in the intensity of the structures of interest. The main contribution of this work

are: definition of a new local-linear region model for the level-set method, and derivation of a computationally efficient approximate flow for the minimization of the level-set energy functional. This model is applied to the problem of automatically segmenting brain tissues in MRI in chapter 3. However, in the context of this thesis, we are also interested in applications related to vascular structure segmentation. Accordingly, work on related level-set segmentation models for vascular structures segmentation have been done, as is briefly outlined in chapter 1 and presented in Appendix II. A limitation of many global segmentation methods, such as the one presented here, is that they are generally sensitive to the presence of adjacent structures with similar intensity characteristics. We argue that the use of a structure-specific local indicator could help to achieve better results in those cases. The definition of such an indicator is the subject of the next section.

2.2.2 Vessel detection and segmentation

Global segmentation methods are generally good at delineating large structures. They are usually less successful in segmenting very fine structures, along which the intensity can vary rapidly, or structures very close to other tissues with similar voxel intensity. One way to overcome this difficulty is to use a highly localized indicator to detect the structures of interest. Devising such a vessel detection filter for the segmentation of fine vascular structures has been defined as the second specific objective of this thesis.

The visualization of various curvilinear structures in 3D CTA or MRI volumes is of the utmost importance in medical imaging. These structures can represent veins, arteries, bronchi, or other structures crucial to physiological functions. By isolating curvilinear structures from the rest of the volume, automatic vessel detection filters can help practitioners visualize the full length of the vessel. These filters are regularly used as a basic building block of many medical image processing applications, such as vessel tracking, segmentation, and image-guided surgery.

With vessel detection filters, the idea is to define functions acting at the pixel level which estimates the degree of membership of a pixel in some geometrical or statistical model. Using sufficiently localized features, it is possible to design functions that are robust to a global change in contrast, or in other words, to design functions that are capable of detecting the desired structures that appear at various intensity levels. Two important qualities of a vessel detection filter are specificity, which is the ability to remove background structures, and sensitivity, which is the ability to detect all parts of the structures of interest. The most used operators in this category are those based on analysis of the Hessian matrix. However, because they implicitly rely on the assumption of on a single vessel per voxel, they tend to suffer from large signal loss at

vessel junctions. Also, the limited information carried by the Hessian matrix is not sufficient to discriminate against all background structures. Consequently, the sensitivity and specificity characteristics of the existing methods are not ideal.

In chapter 4, we define a vessel detection filter based on the analysis of a set of second order differences computed in many directions, which we refer to as a *structure ball*. The definition of this new geometric construction was inspired by recent advances in the processing of high angular resolution diffusion imaging in neurology, which allows for a more accurate tracking of brain tissue fibers. With this structures, it is possible to represent single curvilinear structures as well as X- and Y-bifurcations. Analysis of the shape of the *structure ball* gives us the ability to define vessel membership measures that do not depend on the single vessel assumption. In this research, a new *vesselness* — or curvilinear structure — measure is proposed. A straightforward method is also presented to fix the parameters of the filter, which makes the filter simple to use. Finally, the filter has been applied to scans from two different medical imaging modalities, found to have better sensitivity and specificity characteristics than existing filters.

The main contribution of this work is the definition of a new local structure descriptor, and its use for the detection of curvilinear structures in 3D medical imaging volumes. This part of the thesis methodology enable us to meet our second specific objective, which is to devise a vessel detection filter for the segmentation of fine vascular structures. The proposed method is relatively simple to understand and implement, and is a significant improvement over previous existing techniques. Its potential to produce automatic segmentation of the coronary arteries is also demonstrated. We can speculate that the integration of the proposed method into a more general segmentation framework, such as the level-set method, could make it possible to segment more complex vascular structures.

2.2.3 2D/3D registration for the guidance of surgical procedures

Once a 3D model of the patient structures has been created, it can be used to improve diagnosis and surgical planning. However, with the appropriate computational tool, it could also help with interventional guidance by augmenting the 2D imagery with the 3D geometry segmented from CTA acquisition. A 3D patient specific model aligned and overlaid on the 2D interventional images, constitutes an accurate reference for the surgeon, and greatly reduces the ambiguities inherent in the interpretation of the 2D images. To this end, in the third part of this research, we propose an efficient 2D/3D non-rigid registration methodology for the guidance of complex procedures, which is the third specific objective of this research. Specifically,

we were interested in registering 3D coronary artery models with live fluoroscopy for the guidance of cardiac intervention, with an emphasis on percutaneous coronary interventions (PCI) of chronic total occlusions (CTO).

The guidance of PCI is a subject that is well represented in the literature. Nonetheless, aligning the 3D geometry with the 2D interventional images has proved to be challenging for various reasons: difficulties in making accurate quasi-intermodal correspondences between images of different dimensionalities, differences in the apparent shape of the structures in CTA and fluoroscopy, computational problems when finding the optimal registration transformation, and the high runtime of the algorithms, to name a few. It is possible to classify most approaches that have been proposed for this task in two main categories: image-based and model-based approaches. The former considers the alignment of a substantial portion of the CTA voxels with the interventional images, while the latter makes use of a simplifying model of the structure, generally a centerline model. The image-based approach often involve using measures based on signal theory, such as the normalized cross-correlation, and is, generally speaking, reputed as the most accurate. However, there is a high computational cost associated with the projection of the voxels on the interventional planes, and this cost limits the complexity of the transformation model that can practically be computed. Therefore, most published image-based methods are limited to rigid transformations. As the 3D planning image is generally acquired under a breathhold, there are significant shape changes as compared to the intraoperative images acquired under free breathing. This makes it important to use an affine or a non-rigid registration method to achieve more convincing registrations. In this respect, the lower computational cost associated with model-based approaches is a decisive advantage. Nevertheless, the use of affine and non-rigid transformation models is not well documented in the literature. Furthermore very little work has been published regarding the optimization of the cost-functions used in the registration process.

In chapter 5 we present a new method for the non-rigid 2D/3D registration of coronary artery models with live fluoroscopy for guidance of cardiac interventions. A model-to-model projective approach is proposed, which is composed of two main steps. In the first step, the parameters of a global transformation model — translation only, rigid, or affine — are estimated by minimizing an energy function, and provide the initial alignment. The three transformation models are used in sequence to minimize the risk of finding local minima. The performance of nine different general purpose optimizers to minimize the corresponding energy function have been evaluated. In the second step, a fully non-rigid reconstruction and registration process is defined based on a variational framework. This process allows the 3D centerline model to

snap onto adjacent structures, increasing confidence in the registration process and compensating for deformations caused by respiration and heartbeat. The method was tested on five clinical datasets, including one CTO dataset and another dataset with a significant build-up of plaque. With total runtime typically under 3s for non-rigid registration, the method is reasonably fast and can be used to enhance surgical intervention guidance by displaying a deformed and aligned 3D model alongside the interventional images.

The contribution of the method presented in this section is fourfold: 1) a 2D/3D rigid mono-plane alignment formulation is adapted for affine alignment with biplane datasets, 2) many local and global optimizers are evaluated and compared for the task of estimating the parameters of the various transformation models, 3) a new efficient 2D/3D fully non-rigid registration formulation based on a variational framework is proposed and evaluated on difficult simulated and clinical dataset, and 4) two multiframe formulations are presented, and pave the way toward 3D+t modeling. This multifaceted contribution allows to enhance surgical guidance by using a 3D model of the patient structure, and meets the third specific objective of this thesis. It is important to note that image-based techniques are more appropriate in some situations, for example when it is difficult to create a centerline model of a 3D structure. Work on this aspect of the problem has been carried out, and has been briefly introduced in the literature review (Julien Couet and *et al.*, 2012).

CHAPTER 3

UNSUPERVISED MRI SEGMENTATION OF BRAIN TISSUES USING A LOCAL LINEAR MODEL AND LEVEL SET

David Rivest-Hénault¹ and Mohamed Cheriet¹,

¹ Département de génie de la production automatisée, École de Technologie Supérieure,
1100 Notre-Dame Ouest, Montréal, Québec, Canada H3C 1K3

Published in Elsevier Magnetic Resonance Imaging,
Volume 29, Issue 2, February 2011, Pages 243–259

Abstract

Real world magnetic resonance imaging of the brain is affected by intensity nonuniformity (INU) phenomena which make it difficult to fully automate the segmentation process. This difficult task is accomplished in this work by using a new method with two original features: (1) each brain tissue class is locally modeled using a local linear region representative, which allows us to account for the INU in an implicit way and to more accurately position the region's boundaries; and (2), the region models are embedded in the level set framework, so that the spatial coherence of the segmentation can be controlled in a natural way. Our new method has been tested on the ground-truthed IBSR database and gave promising results, with Tanimoto indexes ranging from 0.61 to 0.79 for the classification of the white matter, and from 0.72 to 0.84 for the gray matter. To our knowledge, this is the first time a region-based level set model has been used to perform the segmentation of real world MRI brain scans with convincing results.

Keywords: IBSR, brain scan, image segmentation, level set, active contours, local linear approximations.

3.1 Introduction

Magnetic resonance imaging (MRI) is a well established and non-invasive means of acquiring high-resolution images of the brain of a live subject. Its precision and good soft tissue contrast permit the discrimination of the nerve connections (white matter or WM) from the congregations of neurons (gray matter or GM) and cerebrospinal fluid (CSF). Analysis of the spatial distribution of those tissues are used to support the diagnosis of various brain illness such

as Alzheimer's disease (Venneri, 2007), dementia with Lewy bodies and Parkinson's disease (Bozzali *et al.*, 2008).

Segmentation of the WM, GM and CSF on a T1-weighted MRI scan is often a prerequisite for such an analysis. As a typical MRI scan involves dozens of high-resolution slices, manual segmentation of the data is tedious and costly. In today's context, where qualified workmanship is scarce and expensive, the automated segmentation of magnetic resonance (MR) images is becoming increasingly relevant. Furthermore, unsupervised methods become efficient tools for automatic segmentation and may have the added benefit of greatly reducing user bias in studies (Vaidyanathan *et al.*, 1997). The major difficulties that arise with the automated segmentation of brain tissues using T1 MR images are related to the fact that image intensities are not constant for each tissue classes. During acquisition, the observed volume is discretized into voxels of significant finite dimension. This leads to a partial volume effect (PVE) since the signal produced by all the tissues present inside a voxel will be averaged, and so the recorded voxel intensities will not, in general, correspond to a single tissue class. In addition, limitations of the imaging device, such as imperfections in radio frequency (RF) coils, result in field inhomogeneities, which can lead to intra-scan intensity nonuniformity (INU) (Mason and Rothman, 2002; Belaroussi *et al.*, 2006). This phenomenon is often referred to as the *bias field* of the MRI system. Finally, inter-scan variability may arise from inappropriate normalization and various other sources of noise may blur the images or displace voxels. Among those sources of uncertainty, the most problematic are PVE and INU. Various difficulties that can occur with MR images are shown in Fig. 3.1.

Although attempts were made to automatically segment MRI brain tissues twenty years ago, the results, in real-world conditions, have only started to become convincing in the last few years. Early segmentation methods were often based on simple clustering techniques and produced discontinuous results with a granulous aspect in the presence of INU and noise. A key aspect of many recent techniques is the notion of *spatial coherence*, which refers to the *a priori* knowledge that structures of interest are continuous in nature. Therefore, recognizing that the probable class of any voxel is related in some way to the classes of its neighbors permits to define powerful methods giving segmentation results of more natural appearance. Approaches which are the most relevant to ours are presented below.

3.1.1 State-of-the-art techniques and their limitations

A great deal of effort has been expended to produce unbiased MR images over the past 15 years. In 1998, Sled *et al.* proposed, in a now classical paper, a retrospective nonparametric method

for INU correction by estimating a slowly varying bias field (Sled *et al.*, 1998). Today, bias field correction remains a very active research topic and more refined methods have emerged (Belaroussi *et al.*, 2006; Manjón *et al.*, 2007). There are, however, situations where such preprocessing techniques have not proved to be beneficial (Ferreira, 2007). Another avenue is the development of prospective methods that attempt to prevent INU artifacts from appearing during the image formation process. Deoni *et al.* presented a method (Deoni *et al.*, 2006) that enables the creation of synthetic scans devoid of RF artifacts. While their results show great potential, this kind of acquisition is not yet part of routine clinical workflow, nor can this method be used for the *a posteriori* correction of an existing dataset. It is therefore desirable to develop segmentation methods that are as resilient as possible to INU.

In 1999, van Leemput *et al.* (van Leemput *et al.*, 1999) published one of the first effective approaches to reliable brain tissues segmentation. They proposed using an EM algorithm to find a four-class partition (WM, GM, CSF and “other”) of observed data. Their method relies on a Markov random field (MRF) for spatial coherence, and is initialized using a template. Marroquin *et al.* (Marroquín *et al.*, 2002) take a similar probabilistic approach. Their main contribution is the definition of a variant of the expectation maximization (EM) algorithm, called MPM-MAP, which is more robust to error in the approximation of the posterior marginal of the MRF. Another related method, the FAST algorithm available in the FMRIB Software Library, is due to Zhang *et al.* (Zhang *et al.*, 2001).

Other authors have used GMMs (Gaussian mixture models) with more than four components in order to account for PVE (Cuadra *et al.*, 2005). Ferreira da Silva (Ferreira, 2007) proposed using a Dirichlet process to find the best N -components GMM and using the N3 tool to compensate for the INU. His method automatically estimates the number of components in the mixture model.

Some disadvantages of the parametric methods are that they all make relatively strong assumptions about the form of the underlying distribution, and they do not cope well with overlapping distributions. Also, imposing a spatial coherence is often an issue (Marroquín *et al.*, 2002). Greenspan *et al.* (Greenspan *et al.*, 2006) introduce a model which takes advantage of a GMM with *many* components to model both the intensity and position of the tissues. Their method is initialized using a K-means procedure.

Non-parametric clustering techniques have also received a great deal of attention for the segmentation of MR brain scans. Siyal and Yu proposed using a modified version of the fuzzy C-means (FCM) (Siyal and Yu, 2005). Their technique, dubbed the intelligent fuzzy C-means

(IFCM), takes advantage of a modified objective function which accounts for the bias field. This technique is very simple and gives good results. Recently, a method that uses a modification of the FCM algorithm inside a three steps process has been introduced by Sikka *et al.* (Sikka *et al.*, 2009) and gives good results. However, spatial consistency is handled in a separated process from segmentation. The classic work of Pham and Prince also benefits from the fuzzy framework (Pham and Prince, 1999, 2004).

Bazin and Pham (Bazin and Pham, 2007) embedded the FCM algorithm in an interesting topology-preserving framework. Their model is initialized with a template which defines the brain topology. This template then evolves, in a topology-preserving fashion, by iterating steps of thinning, class membership evaluation and dilation. The main drawback of this technique is the cost of a good topological template: manual editing may be required. Also, this method will not give good results if the topology of the brain scan does not fit the template well.

An alternative non-parametric approach is presented in (Jimenez-Alaniz *et al.*, 2006). The authors start by clustering the data using a version of the mean shift (MS) algorithm (Comaniciu and Meer, 2002). Those clusters are then classified with the help of probability templates. However, no mechanism is provided to counterbalance the INU effect. A related, but different MS-based procedure which does not rely on any template is presented in (Mayer and Greenspan, 2009). From another perspective, wavelet-based segmentation technique, such as the one proposed by Zhou and Ruan (Zhou and Ruan, 2007) are starting to emerge.

Variational segmentation methods using active contours have gained momentum in the last few years. Chan and Vese's *Active Contour Without Edges* was the first variational region-based model to become popular. They provide a model for the piecewise constant and the piecewise smooth cases (Vese and Chan, 2002). Subsequently, interesting segmentation models have been proposed by many researchers (Tsai *et al.*, 2001; Gao and Bui, 2005; Paragios, 2006).

Very few authors have used level set based models to segment MR brain images, although some examples can be found (Gao and Bui, 2005; Vese and Chan, 2002; Angelini *et al.*, 2007). Unfortunately, INU is not well handled by their constant tissues model and they do not provide detailed results on the well known Brainweb or IBSR dataset.

Recently, some authors have introduced level set methods for brain tissues segmentation which benefit from localized region models and presented impressive results (Li *et al.*, 2005b, 2008b,a). Nevertheless, the performance of level set-based methods for brain MRI segmentation is still scarcely documented, and those methods would benefit from more thorough evaluations on

well known datasets. Moreover, the flexibility of the active contour variational framework allows for the definition of other, and hopefully even more accurate, segmentation models.




Figure 3.1 Two consecutive sample slices from the IBSR, subject 5_8: (a) slice no.15, and (b) slice no.16, enhanced for visualization purposes. Area A and area B in (a) are of about the same intensity, even though they represent different classes (GM and CSF). Area A is dark because of the PVE. (b) shows a low frequency wave effect over the brain domain, which can be attributed to the bias field. Because of this, area C is significantly lighter than the surrounding area.

3.1.2 The need for a more comprehensive model

From the above discussion, it appears that the spatial coherence of the results is still an issue. Active contour schemes allow for the inclusion of a regularization parameter which controls the viscosity of the contour and which may lead to more natural results. Of these tools, the level set method is one of the best available, for the following reasons: its generality allows a large variety of energy terms to be represented; it automatically handles topological changes; and it is relatively simple to implement.

Many segmentation methods rely on some kind of template for the inclusion of *prior* knowledge. That knowledge may be misleading when the brain scan shows an atypical structure, such as with infant brain data or a brain with tumor data. Also, the production of an appropriate template may be expensive. Given equivalent performance, the method which uses the simplest template will generally be considered superior.

It seems clear that INU must be taken into account, since all the most successful methods include a way of modeling the bias field. However, all the methods reviewed use an explicit model for the bias field. This adds to the model complexity, and may conflict with the rest of the segmentation process.

In this paper, we present a new, fully unsupervised algorithm for the segmentation of MR brain images. This new algorithm takes advantage of the level set method to control the spatial coherence of the segmentation. An original local linear region representative is used to model various classes. This type of region representative, introduced in (Rivest-Hénault and Cheriet, 2007) and based on the work of (Brox, 2005; Liu, 2006; Li *et al.*, 2008b), makes it possible to closely follow the low frequency INU within a region. Our method offers several advantages over other classical level set methods, such as the piecewise smooth model of Chan and Vese (Vese and Chan, 2002): 1) Not only is fitness of the voxel intensities tested by our functional, but the fitness of the voxel intensities spatial derivatives are tested as well. This is a valuable segmentation cue for images with significant gradient. In this respect, a level set method may require fewer iterations to reach its final position; 2) As the model is defined locally for each node, a fast narrow band implementation is possible; and 3) As our model has a local statistical definition, it is not necessary to use another model to define the region's representatives outside their original domain.

The method presented in this paper is related to a few recent state-of-the-art methods (Rosenhahn *et al.*, 2007; Brox and Cremers, 2009; Li *et al.*, 2008b; Brox *et al.*, 2010). Among them, the work of Li *et al.* (Li *et al.*, 2008b) is perhaps the closest. However, there are important distinctions: 1) While Li *et al.* presented a generic segmentation tool, we have focused on the definition of a method for brain tissue segmentation in particular. This has implications for the initialization procedure, as well as for the treatment of the outliers; 2) Li *et al.* proposed a purely functional formulation of the segmentation problem, whereas ours is probabilistically motivated; 3) Our method benefits from a local linear region model, different from that of Li *et al.*; 4) Finally, we present results obtained on both simulated and real-world, well-established datasets.

3.1.3 Organization of this paper

Section 3.2 gives details of our new method. Section 3.3 presents experimental results on synthetic and natural images, and, finally, section 3.4 discusses the proposed model.

3.2 A local linear level set approach to MRI segmentation

3.2.1 Formulation of the problem

Let $\Omega \in \mathbb{R}$ be the image domain, $R \subset \Omega$ a certain image region and $P(I | R)$ the assumed distribution model of the pixels in a region R . It is assumed that all the image pixels on the domain

$\Omega \in \mathbb{R}$ are independent and that the observed noise is random. The maximum *a posteriori* segmentation of the image then consists in the set $\{\hat{R}_k\}_{k=1}^N$ of N regions that maximizes the problem:

$$\{\hat{R}_k\}_{k=1}^N = \arg \max_{R_k \in \Omega} P(\{R_k\}_{k=1}^N | I)$$

or

$$\{\hat{R}_k\}_{k=1}^N = \arg \max_{R_k \in \Omega} P(I | \{R_k\}_{k=1}^N) P(\{R_k\}_{k=1}^N)$$

using Bayes' rule (Ben Ayed *et al.*, 2006). The above problem can be converted into a negative log-likelihood minimization problem, as follows:

$$\{\hat{R}_k\}_{k=1}^N = \arg \min_{R_k \in \Omega} E[\{R_k\}_{k=1}^N]$$

where the energy $E[\{R_k\}_{k=1}^N]$ is given by:

$$E[\{\hat{R}_k\}_{k=1}^N] = \sum_{k=1}^N \left[\int_{\mathbf{x} \in R_k} -\log P(I(\mathbf{x}) | R_k) d\mathbf{x} - \log P(R_k) \right]. \quad (3.1)$$

This energy now depends on suitable probability functions $P(I | R)$ and $P(R)$.

3.2.2 Local linear region model

Here, we propose a new segmentation model that uses linear approximations as the region representatives, and which shares some similarities with the work by Bui *et al.* (Bui *et al.*, 2005). However, as pioneered by (Liu, 2006; Mory *et al.*, 2007; Li *et al.*, 2008b; Brox and Cremers, 2009), we choose to compute a local approximation of the model for each pixel position, instead of computing one global model.

Using the above probabilistic notation and neglecting the normalization term, we introduce the following probability of a given pixel in a certain image region:

$$P(I | R) = \exp \left(-\alpha \frac{(I - [a + bx + cy])^2}{\sigma_e^2} \right) \cdot \exp \left(-\mu \left[\frac{(I_x - b)^2}{\sigma_{eb}^2} + \frac{(I_y - c)^2}{\sigma_{ec}^2} \right] \right), \quad (3.2)$$

which assumes a Gaussian distribution of the error (Gudbjartsson and Patz, 1995). Here, I is an image pixel located at coordinates $\{x, y\}$, I_x and I_y are its first derivative in the x and y directions, a , b and c are the parameters of the linear model ($a + bx + cy = 0$) used as the local region representative, and $\alpha > 0$ and $\mu \geq 0$ are free parameters used to balance the terms. The various σ values are defined as the standard deviation of the error of each term.

A practical concern which appears when using (3.2) for image segmentation is that the edges induce gradients that are not part of any image object. This is not desirable, since the edges I_x and I_y are no longer representative of the region. A straightforward way to handle this situation is to use another function to relax the gradient term when the edge probability is high. The resulting probability density function is as follows:

$$P(I | R) = \exp \left(-\alpha \frac{(I - [a + bx + cy])^2}{\sigma_e^2} \right) \cdot \exp \left(-\mu \cdot g(I) \left[\frac{(I_x - b)^2}{\sigma_{eb}^2} + \frac{(I_y - c)^2}{\sigma_{ec}^2} \right] \right). \quad (3.3)$$

where $g(I)$ is a monotonically decreasing function of the source image gradient, as in the *Geodesic Active Contour* model (Caselles *et al.*, 1997). In this paper, we use the function $g(I) = \frac{1}{1 + |\nabla I|}$, suggested by (Kass *et al.*, 1988). The relaxation function $g(I)$ was absent of the model presented in (Rivest-Hénault and Cheriet, 2007) leading to computation of irrelevant region representatives around salient edges.

The energy (3.1) is hard to minimize when the σ terms are not constant. In order to circumvent this problem, a strong, but generally accepted assumption is that the errors have the same standard deviation everywhere. Based on that assumption, the maximum *a posteriori* energy can be computed with

$$P(I | R) = e^{-\alpha(I - [a + bx + cy])^2} \cdot e^{-\mu g(I) [(I_x - b)^2 + (I_y - c)^2]}. \quad (3.4)$$

Regional parameters a , b , and c are unknown and to be estimated from the data *at each voxel location*.

A local estimate of those parameters can be computed by minimizing the following negative log-likelihood function for a region defined by a windowing function W and centered at a voxel location \mathbf{x}_0 :

$$E_W \left[\{\hat{R}_k\}_{k=1}^N \right] = \int_{\mathbf{x} \in R_k} -\log P(I(\mathbf{x}) | R_k) W(\mathbf{x}) d\mathbf{x} \quad (3.5)$$

with

$$W(\mathbf{x}) = \begin{cases} 1, & \text{if } \max_d |\mathbf{x} - \mathbf{x}_0| < w \\ 0, & \text{otherwise} \end{cases}$$

where d is a dimensional index and w is the half-window size. This problem has a unique solution in most practical situations. The resolution of (3.5) is covered in the next section.

3.2.3 Segmentation using the level set method

The level set formalism (Osher and Sethian, 1988) has been used to implement our segmentation method. Within this framework, a domain is partitioned in up to two phases using a level set function, $\phi(\mathbf{x})$. By definition, $\phi(\mathbf{x}) > 0$ over $\omega \subseteq \Omega$, $\phi(\mathbf{x}) < 0$ over $\Omega \setminus \omega$ and $\phi(\mathbf{x}) = 0$ on $\partial\omega$, the boundary between ω and $\Omega \setminus \omega$. A partitioning curve C is then represented by $\partial\omega$ and its position can be modified by updating $\phi(\mathbf{x})$. The rest of this section describes a two-phase implementation of our local linear segmentation algorithm within the level set framework.

Within the level set framework, the local energy function (3.5) with $P(I(\mathbf{x}) | R_k), k \in \{-, +\}$ defined by (3.4) can be reformulated as follows:

$$\begin{aligned} F_W(\mathbf{x}, a_i, b_i, c_i, \phi) = & \\ & \alpha \int_W (a_- + b_-x + c_-y - u_0)^2 (1 - H(\phi)) d\mathbf{x} \\ & + \alpha \int_W (a_+ + b_+x + c_+y - u_0)^2 H(\phi) d\mathbf{x} \\ & + \mu \int_W g(u_0) [(b_- - \frac{\partial u_0}{\partial x})^2 + (c_- - \frac{\partial u_0}{\partial y})^2] (1 - H(\phi)) d\mathbf{x} \\ & + \mu \int_W g(u_0) [(b_+ - \frac{\partial u_0}{\partial x})^2 + (c_+ - \frac{\partial u_0}{\partial y})^2] H(\phi) d\mathbf{x}, \end{aligned} \quad (3.6)$$

where $u_0 = I$ (the input image), $H(\phi)$ is the Heaviside step function and $i = \{+, -\}$. The indices + and - denote variables which apply to the phase $\phi(\mathbf{x}) > 0$ and $\phi(\mathbf{x}) < 0$ respectively. To compute a_+ , b_+ , and c_+ , we optimize the parameter of this functional so that

$$\frac{\partial F_W}{\partial a_+} = 0, \quad \frac{\partial F_W}{\partial b_+} = 0, \quad \frac{\partial F_W}{\partial c_+} = 0. \quad (3.7)$$

Omitting the ϕ parameter of the $H(\phi)$ function, (3.7) expands to the following system of linear equations:

$$\begin{aligned} a_+ \int_W H d\mathbf{x} + b_+ \int_W xH d\mathbf{x} + c_+ \int_W yH d\mathbf{x} \\ = \int_W u_0 H d\mathbf{x} \end{aligned} \quad (3.8a)$$

$$\begin{aligned} a_+ \int_W xH d\mathbf{x} + b_+ \int_W (x^2 + \frac{\mu g(u_0)}{\alpha}) H d\mathbf{x} \\ + c_+ \int_W xyH d\mathbf{x} = \int_W (xu_0 + \frac{\mu g(u_0)}{\alpha} \frac{\partial u_0}{\partial x}) H d\mathbf{x} \end{aligned} \quad (3.8b)$$

$$\begin{aligned} a_+ \int_W yH d\mathbf{x} + b_+ \int_W xyH d\mathbf{x} + \\ c_+ \int_W (y^2 + \frac{\mu g(u_0)}{\alpha}) H d\mathbf{x} = \int_W (yu_0 + \frac{\mu g(u_0)}{\alpha} \frac{\partial u_0}{\partial y}) H d\mathbf{x}, \end{aligned} \quad (3.8c)$$

or, using matrix notation,

$$A \cdot \theta = B, \quad (3.9)$$

where $\theta = (a_+, b_+, c_+)^T$, $A =$

$$\begin{pmatrix} \sum_W H & \sum_W xH & \sum_W yH \\ \sum_W xH & \sum_W (x^2 + \frac{\mu g(u_0)}{\alpha}) H & \sum_W xyH \\ \sum_W yH & \sum_W xyH & \sum_W (y^2 + \frac{\mu g(u_0)}{\alpha}) H \end{pmatrix},$$

and

$$B = \begin{pmatrix} \sum_W u_0 H(\phi) \\ \sum_W (xu_0 + \frac{\mu g(u_0)}{\alpha} \frac{\partial u_0}{\partial x}) H(\phi) \\ \sum_W (yu_0 + \frac{\mu g(u_0)}{\alpha} \frac{\partial u_0}{\partial y}) H(\phi) \end{pmatrix}.$$

The expressions for a_- , b_- , c_- can be deduced from the above, but using $(1 - H(\phi))$ instead of $H(\phi)$. It can be of interest to note that, when $\mu = 0$, (3.9) is equivalent to the least squares fitting of a plane using the Moore-Penrose pseudoinverse.

The system of linear equations (3.9) is solvable in most practical situations, however special care must be taken to handle the few limiting cases where the system becomes under-determined. In all cases, reasonable assumptions can be made to set the value of a_+ , b_+ , and c_+ . Let us denote by N_W the number of voxels where $H > 0$ inside a window W . The limiting cases are as follows: 1) $N_W = 0$, then $\sum_W H = 0$, a valid solution being $\theta = \mathbf{0}$. 2) $N_W = 1$, a single voxel can be assimilated to a plane with parameters $a_+ = u_0$, $b_+ = 0$ and $c_+ = 0$. 3) $N_W \geq 2$, but the y values are constant, i.e. all voxels lie on a line parallel to the x axis, a solution being to fix $b_+ = 0$ and to ignore (3.8b). 4) The same as case 3, but replacing x by y and b_+ by c_+ and (3.8b) by (3.8c). 5) All voxels lie on a line passing through the origin, i.e. $y_i = kx_i$ with $k \in \mathbb{R}$. This is very unlikely to occur numerically because of the definition of the grid, and has been ignored in our implementation. However, it could be possible to add the condition $c_+ = kb_+$ to the linear system. Case 1 is of limited interest, since it corresponds to the situation where no valid voxel lies inside the window W , and, thus, the window W is centered in a location where the level set function is always negative and far away from the segmentation boundaries. Cases 2–5 can be handled by careful programming. In all other situations, (3.9) is solvable using classical techniques. It is also possible to derive explicit expressions for parameters a_i , b_i and c_i . Therefore, they are fully defined and their values depend only on $\phi(\mathbf{x})$ and on the input image.

We now introduce a functional which can be used for a two-phase segmentation of an image with the level set method, as in (Chan and Vese, 2001). Using equations (3.1) and (3.4), our

new model's functional is written as follows:

$$\begin{aligned}
F(a_i, b_i, c_i, \phi) = & \\
& \alpha \int (a_- + b_-x + c_-y - u_0)^2 (1 - H(\phi)) d\mathbf{x} \\
& + \alpha \int (a_+ + b_+x + c_+y - u_0)^2 H(\phi) d\mathbf{x} \\
& + \mu \int g(u_0) [(b_- - \frac{\partial u_0}{\partial x})^2 + (c_- - \frac{\partial u_0}{\partial y})^2] (1 - H(\phi)) d\mathbf{x} \\
& + \mu \int g(u_0) [(b_+ - \frac{\partial u_0}{\partial x})^2 + (c_+ - \frac{\partial u_0}{\partial y})^2] H(\phi) d\mathbf{x} \\
& + \nu \int |\nabla H(\phi)| d\mathbf{x}.
\end{aligned} \tag{3.10}$$

This functional is composed of three terms. The first is for the error between the linear approximation and the source image. The second measures the difference in the gradient between our region approximation and that of the source image. The third term accounts for the description length and is a regulator. The regional parameters a_i , b_i and c_i vary in function of the position and are computed as previously described. Their values depend on ϕ , and therefore, should be considered as embedded in the above functional.

Starting from some initial position, the segmentation boundaries are to be optimized up to a local minima using an approximate gradient descent method. The level set evolution equation is written as

$$\begin{aligned}
\frac{\partial \phi}{\partial t} = & \delta(\phi) \left[\nu \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) + \alpha (a_- + b_-x + c_-y - u_0)^2 \right. \\
& + g \cdot \mu \left[(b_- - \frac{\partial u_0}{\partial x})^2 + (c_- - \frac{\partial u_0}{\partial y})^2 \right] \\
& - \alpha (a_+ + b_+x + c_+y - u_0)^2 \\
& \left. - g \cdot \mu \left[(b_+ - \frac{\partial u_0}{\partial x})^2 + (c_+ - \frac{\partial u_0}{\partial y})^2 \right] \right].
\end{aligned} \tag{3.11}$$

This equation is derived from the partial derivative of (3.10) by respect to ϕ , but neglecting small terms that are proportional to the ratio between the length of the contour and the area of the region. Similar approaches have been proposed by (Brox, 2005; Rosenhahn *et al.*, 2007). For a more detailed description of the approximation, the reader is referred to the Appendix. Since, the values of the regional parameters a_i , b_i and c_i depend on ϕ , they need to be updated at each iteration. This can be done directly by replacing these terms by their explicit formulation, or by solving (3.9) at each iteration.

When compared to the smooth *Active Contour Without Edge* model (Chan and Vese, 2001), it is expected that our new local linear segmentation model will give more accurate results when applied to the segmentation of an image with a significant gradient. For a simple example, see Fig. 3.2.




Figure 3.2 Our local linear model **B** with its local linear approximations, successfully segmented the shape. The piecewise smooth model **A** was not able to discern the darkest part of the shape from the background.

3.2.4 Extension to a 4 phase model

As the brain tissues are to be classified in three classes (WM, GM and CSF), the model has to be extended to the 3-phase case, at a minimum. The approach presented in (Vese and Chan, 2002) takes advantage of coupled curve evolution and allows for the representation of up to 2^p phases using p level set functions. In the 4-phase case, two LSFs evolve together. Let the functions u_{++} , u_{+-} , u_{-+} , u_{--} , constitute the restriction of u on the four phases. They are defined as follows:

$$u(\mathbf{x}) = \begin{cases} u_{++}(\mathbf{x}), & \text{if } \phi_1(\mathbf{x}) > 0 \text{ and } \phi_2(\mathbf{x}) > 0 \\ u_{+-}(\mathbf{x}), & \text{if } \phi_1(\mathbf{x}) > 0 \text{ and } \phi_2(\mathbf{x}) < 0 \\ u_{-+}(\mathbf{x}), & \text{if } \phi_1(\mathbf{x}) < 0 \text{ and } \phi_2(\mathbf{x}) > 0 \\ u_{--}(\mathbf{x}), & \text{if } \phi_1(\mathbf{x}) < 0 \text{ and } \phi_2(\mathbf{x}) < 0 \end{cases} \quad (3.12)$$

In our experiments, we take advantage of 3 or 4 phases, depending on the data. When only three phases are needed, one of them, u_{--} , will remain empty. With this notation, the level set

evolution function for ϕ_1 is described by the equation

$$\begin{aligned}
\frac{\partial \phi_1(\mathbf{x})}{\partial t} = & \delta(\phi_1) \left[\nu \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \right. \\
& - H(\phi_2) \cdot \left(\alpha [(a_{++} + b_{++}x + c_{++}y - u_0)^2 \right. \\
& \quad \left. - (a_{-+} + b_{-+}x + c_{-+}y - u_0)^2] \right. \\
& \quad \left. + g \cdot \mu \left[(b_{++} - \frac{\partial u_0}{\partial x})^2 + (c_{++} - \frac{\partial u_0}{\partial y})^2 \right. \right. \\
& \quad \left. \left. - (b_{-+} - \frac{\partial u_0}{\partial x})^2 - (c_{-+} - \frac{\partial u_0}{\partial y})^2 \right] \right) \\
& - (1 - H(\phi_2)) \cdot \left(\alpha [(a_{+-} + b_{+-}x + c_{+-}y - u_0)^2 \right. \\
& \quad \left. - (a_{--} + b_{--}x + c_{--}y - u_0)^2] \right. \\
& \quad \left. + g \cdot \mu \left[(b_{+-} - \frac{\partial u_0}{\partial x})^2 + (c_{+-} - \frac{\partial u_0}{\partial y})^2 \right. \right. \\
& \quad \left. \left. - (b_{--} - \frac{\partial u_0}{\partial x})^2 - (c_{--} - \frac{\partial u_0}{\partial y})^2 \right] \right) \left. \right] \quad (3.13)
\end{aligned}$$

It is possible to deduce the equation for $\frac{\partial \phi_2(\mathbf{x})}{\partial t}$ by swapping ϕ_1 and ϕ_2 and by rearranging the signs of the region errors.

3.2.5 Outliers rejection

Whenever an isolated region shrinks until it becomes sufficiently small, the error terms of (3.11) approach zero for this region, because the statistics are computed locally. This prevents this region from collapsing on itself, even if it shows very little contrast with its surrounding. It can be argued that in the case where few samples are available, the local statistics are not good representatives of regions. To circumvent this problem, we introduce a simple form of outliers rejection: whenever the number of sample $n < W_{min}$ inside a window W , the errors terms of (3.11), $(a_i + b_i x + c_i y - u_0)^2$, $(b_i - \frac{\partial u_0}{\partial x})^2$ and $(c_i - \frac{\partial u_0}{\partial y})^2$ are replaced by a large constant K for this region. In this work, we used $K = \max(u_0)^2 = 255^2$. This has the effect of making the region collapse on itself and letting the surrounding regions compete for the available space.

3.2.6 Implementation issues

In practice, the Heaviside and Dirac functions of the curve evolution equations must be regularized. For example, a common choice (Li *et al.*, 2006) of a finite support regularization

is

$$\begin{aligned}
 H_\varepsilon(z) &= \begin{cases} 0 & \text{if } z < -\varepsilon \\ 1 & \text{if } z > \varepsilon \\ \frac{1}{2}(1 + \frac{z}{\varepsilon} + \frac{1}{\pi} \sin(\frac{\pi z}{\varepsilon})) & \text{otherwise} \end{cases} \\
 \delta_\varepsilon(z) &= \begin{cases} 0 & \text{if } \text{abs}(z) > \varepsilon \\ \frac{1}{2\varepsilon}(1 + \cos(\frac{\pi z}{\varepsilon})) & \text{otherwise} \end{cases}
 \end{aligned} \tag{3.14}$$

where ε is the regularization parameter. In this paper, we chose to use only a slightly regularized version of the Heaviside function, and we fixed $\varepsilon = 1$ for all our experiments.

The use of local operations in (3.8) is computationally expensive. However, as it was noticed in (Tsai *et al.*, 2001), the values for a, b and c change slowly through time, so they need not to be computed at every iteration. In this work, we update them every $25 \approx w/3$ iterations (see sec 3.2.10), which works well in practice: the results obtained this way are very similar to those obtained while updating the a, b and c at every iteration.

3.2.7 3D extension of the model

As our model is based on the level set methodology, it is easily expandable to multidimensional cases. Two possibilities are described here. The first is to use 3D surfaces to represent the segmentation boundaries, but to continue to compute 2D local linear region representatives for each slice. This is potentially useful for volumes where the mean intensity of the bias field varies greatly from one slice to another. A 3D surface can be represented using a 3D level set function, without much added difficulties (Sethian, 1999) and will ensure a certain degree of 3D smoothness.

A second possibility is again to use 3D surfaces to represent the segmentation boundaries, but to use 3D instead of 2D local linear models. If the volume is affected by a single bias field, this would exploit the 3D nature of the data better. In this case, the most significant change is the use of a multidimensional linear approximation, with $d + 1$ components, for the regions' representatives in the functionals (3.10) and (3.6). Equation (3.3) is replaced by

$$\begin{aligned}
 P(I | R) &= \exp \left(-\alpha \frac{(I - [a + bx + cy + dz])^2}{\sigma_e^2} \right) \\
 &\quad \cdot \exp \left(-\mu \cdot g(I) \left[\frac{(I_x - b)^2}{\sigma_{eb}^2} + \frac{(I_y - c)^2}{\sigma_{ec}^2} + \frac{(I_z - d)^2}{\sigma_{ed}^2} \right] \right),
 \end{aligned}$$

and (3.10)–(3.6) are adjusted accordingly. The region representatives are actually computed in the same way as in (3.7) and the LSF evolution is a natural extension of (3.11).

In both cases, if the data are anisotropic, special care should be taken when computing the derivative of (3.11) numerically. We think that the use of a 3D implementation is to be recommended because that enforces an inter-slice spatial coherence. However, in this paper, we have chosen a slice-by-slice implementation because of time constraints.

3.2.8 Computational cost

All the operations in (3.11) are linear in the number of pixels, n , and are then performed in time $O(n)$. The integrals of (3.8) at each pixel location are evaluated in time $O(n \cdot m^2)$ where m is the size of the window. Other operations of (3.8) are linear in n . Overall, the computational time for our algorithm is then $O(n) + O(n \cdot m^2) + O(k) = O(n \cdot m^2)$ per iteration. However, as there is no global image operation, a good acceleration is expected from a narrow band implementation. For a typical $256 \times 256 \times 60$ IBSR volume processed with a window size of $w = 71$, our running times on an Intel core 2.00GHz Ubuntu Linux laptop are around 40s for the FCM initialization and about 30 minutes per volume for our algorithm implemented in Matlab. If the window size is increased to $w = 91$, the running time of our algorithm increases to 33 minutes.

3.2.9 Initialization by Fuzzy C-means clustering

Level set segmentation models which use smooth region representatives are well known to fall into a local minimum if initialized far from the desired solution. In order to converge to a valuable local minimum, we propose to use the classical Fuzzy C-means (FCM) clustering algorithm with 3 clusters to provide an initialization point for our segmentation model. In this paper, we fixed the FCM's fuzziness parameter to 3 and we use the Matlab implementation of the algorithm.

The labels associated with the three FCM clusters have to be set carefully, as a function of the data. When segmenting a simulated brain scan, we used WM, GM, and CSF. However, when segmenting a real-world brain scan, we used WM, GM, and DG (dark gray). An additional partition which represents CSF is then created by thresholding: the 5% darker voxels are labeled CSF. This procedure can be motivated by the data. For the simulated scans, the number of voxels forming the WM, GM, and CSF classes are comparable. By contrast, for the real-world IBSR 20 normal brain dataset, the CSF class only accounts for 0.9–2.1% of the total number

of voxels, and such a small representation is not well handled by the FCM algorithm. Also, we found that the voxel intensity of the GM class in a real-world scan expresses as much as twice the variance as that of the WM class. This encouraged us to use two clusters, GM and DG, to represent the gray matter when segmenting real-world scans. Please refer to sec. 3.3 for more details on the datasets.

Different standalone versions of the FCM algorithm were used for the segmentation of MR brain data with some success. The principal difficulties in the approach seem to come from a lack of spatial coherence. This is a particular aspect that we want to address with our level set model.

3.2.10 Overview of our new segmentation model

The principal steps of our algorithm are set out in Algorithm 3.1. It is worth noting that a fully 3D implementation is used for the FCM initialization. In contrast, the local level set model was implemented in slice-by-slice fashion, each slice being processed independently.

Algorithm 3.1 Automatic segmentation of brain tissues

- 1: Load the MR brain scan data into SCAN.
 - 2: Partition SCAN using the FCM algorithm, store the result in CLUSTERS.
 - 3: Label each voxel according to the class of CLUSTERS for which its FCM membership value is the highest.
 - 4: Label the darker voxels as CSF (if needed, see sec. 3.2.9).
 - 5: Set PHI1 and PHI2 according to (3.12).
 - 6: Perform region optimization with our local linear model by evolving PHI1 and PHI2 with (3.13). This iterative procedure stops when no point on PHI1 or PHI2 changes sign during an iteration.
-

The value of the various parameters for our model have to be handpicked by trial and error on test images. However, once the model parameters have been adjusted for a specific task, the segmentation results are robust to a small perturbation of these values. Therefore, most of the parameters were fixed for *all* the experiments in section 3.3. These values are: $\alpha = 1$, $\mu = 50$, $\nu = 1 \times 10^{-3} \times 255^2$ and $W_{min} = \frac{5}{3}w$. We adjusted the parameter w in function of the database. Its value was set to $w = 71$ in all our experiments except the last one where we used $w = 91$. The behavior of our method when the value of some parameters change is illustrated in Fig. 3.3.

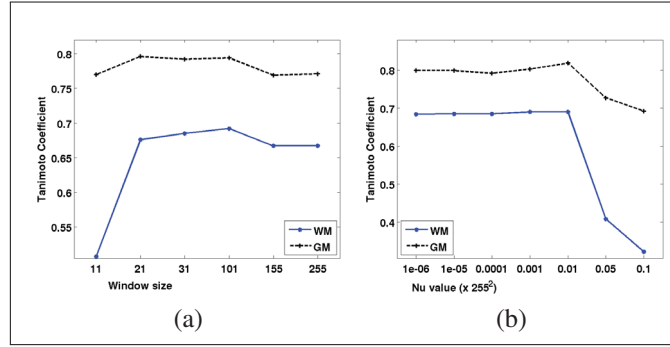


Figure 3.3 Results obtained on a slice of IBSR subject 4_8 keeping all parameters fixed, and changing: (a) the window size; and (b) the ν value. Details for the Tanimoto coefficient are presented in Sec. 3.3.2.

3.3 Experimental validation

3.3.1 Synthetic Data

The BrainWeb MR simulator (Collins *et al.*, 1998) was used to generate synthetic MR brain scans with a level of noise ranging from 0% to 9% and with a INU level of 0%, 20% or 40%. Those data were obtained from the BrainWeb Database at the McConnell Brain Imaging Centre of the Montreal Neurological Institute, McGill University. The synthetic scans used in our experiments were the 8-bit $181 \times 217 \times 60$ voxel brain-only data.

In our experiments, we have used scans with slices 1.0mm thick for comparison purpose with (van Leemput *et al.*, 1999; Marroquín *et al.*, 2002). We also have used slices 3.0mm thick, because they are more similar to the IBSR real data that we also used. The 3.0mm slices are significantly harder to segment, because of an increased PVE which blurs the voxels.

In the literature, the preferred performance index computed for this database is the Dice index, which is defined as $D(k) = \frac{2 \cdot V_{p \cap g}(k)}{V_p(k) + V_g(k)}$, where $V_{p \cap g}(k)$ is the number of pixels classified as class k by both the proposed method and ground-truth segmentation, and $V_p(k)$ and $V_g(k)$ are the number of pixels classified as k by either the proposed method or the ground-truth method.

Qualitative assessment of the results indicates that our method is capable of giving very good results in most situations, even in the presence of a large INU. See Fig 3.4 for a sample result. The global performance indices computed on our results on the BrainWeb database are shown in Fig 3.5 and 3.6 and are compared with the standard FCM algorithm that we used as our initialization.

For the 1mm datasets, Fig. 3.5, it can be seen that, our model compares well with the others ((van Leemput *et al.*, 1999; Marroquín *et al.*, 2002)), especially for high level of noise and INU. As for the 3mm datasets, Fig. 3.6 clearly show that our method greatly improves the precision of the segmentation over the FCM initialization when the INU level is high. This is important because typical real MR brain scans tend to present high level of INU while the presence of other kind of noise remains low. The comparison with the FCM algorithm, used for initialization, demonstrates the improvement of the results due to the level set optimization.

3.3.2 Real Data

The ground-truthed *Internet Brain Segmentation Repository* (IBSR) database *20 normal MR brain data sets and their manual segmentation*¹ presents various levels of difficulty, is freely available online and has previously been used many times in published studies (Worth, 1996). The ground-truth segmentation was performed by a trained investigator using the semi-automated intensity contour mapping algorithm described in (Kennedy *et al.*, 1989). From the available data format, the 8-bit $256 \times 256 \times 64$ voxel brain-only data files were used in our experiments. Slice are 3.0mm or 3.1mm thick, depending on the subject.

We tested our new segmentation model on the 20 subjects in the database using the same set of parameters that was selected during the experiment on the synthetic dataset, which is a indication of the robustness of our algorithm. For those scans, however, we used a four-class segmentation model, as discussed in sec. 3.2.9. At the end of the procedure, voxels in the DG class were converted to the GM class. Prior to the application of our algorithm, we applied a simple normalization algorithm (Johnston *et al.*, 1996), as it is common practice with real data (Mayer and Greenspan, 2009).

The standard performance index used with the IBSR is the Tanimoto coefficient, defined as $T(k) = \frac{V_{p \cap g}(k)}{V_{p \cup g}(k)}$, where $V_{p \cap g}(k)$ is the number of pixels classified as class k by both the proposed method and ground-truth segmentation, and $V_{p \cup g}(k)$ is the number of pixels classified as k by either methods. Note that, even if both performance indices range from 0 to 1, $T(k) < D(k)$, save for the extreme cases.

Figure 3.7 shows a sample of our results. Again, a qualitative assessment indicates that our method leads to a satisfactory segmentation. The performance indices obtained for each subject

¹Provided by the Center for Morphometric Analysis at Massachusetts General Hospital, available at <http://www.cma.mgh.harvard.edu/ibsr/>.

Table 3.1 Average Tanimoto index for various segmentation methods on the IBSR 20 normal brain scan, based on published data. FCM+t is our initialization. The dagger (†) indicates works that do not provide individual results for this dataset.

Method	WM	GM	CSF
Akselrod-Ballin <i>et al.</i> (Akselrod-Ballin <i>et al.</i> , 2006) [†]	0.669	0.680	0.346
Tissue tracking (Melonakos <i>et al.</i> , 2007) [†]	0.658	0.624	-
Ferreira da Silva (Ferreira, 2007) [†]	0.735	0.675	-
IFCM (Siyal and Yu, 2005) [†]	0.724	0.750	-
MPM-MAP (Marroquín <i>et al.</i> , 2002)	0.682	0.657	0.227
Dual-front (Li <i>et al.</i> , 2005b)	0.670	0.739	-
FMRIB FAST (Zhang <i>et al.</i> , 2001)	0.648	0.556	-
MFC (Pham and Prince, 2004)	0.712	0.586	-
MS-edge (Jimenez-Alaniz <i>et al.</i> , 2006)	0.628	0.594	0.210
AMS (Mayer and Greenspan, 2009)	0.688	0.686	0.140
CGMM (Greenspan <i>et al.</i> , 2006) [†]	0.660	0.680	-
FCM (Worth, 1996)	0.567	0.473	0.058
Manual (Worth, 1996)	0.832	0.876	-
FCM+t	0.701	0.753	0.176
Local Linear	0.726	0.795	0.263

are summarized in Fig. 3.8, and average results are shown in Table 3.1². Results labeled FCM+t correspond to our initialization method. Our average and individual results clearly indicate that our method expresses less variance than the others, while at the same time they are more precise.

By comparing the curves *FCM+t* and *Local-Linear* in Fig. 3.8, it can be seen that our local linear model makes it possible to improve on our initialization on difficult subjects, e.g. 16_3, by up to 0.07 points for both the WM and the GM. However, on easier subjects, e.g. 205_3, the benefit of our local linear model is less significant for the WM: our initialization procedure already provides relatively good results. This suggests that the INU effect is not as strong for those volumes. Nevertheless, the application of our method resulted in an appreciable improvement of the segmentation score for the GM with all subjects.

Our results show that most of the CSF regions located in the internal part of the brain are mainly under-segmented (see Fig. 3.7 for some examples). This can be explained by the very large variance expressed by the CSF regions, which makes their segmentation very challenging.

²Please take note that the average result for Ferreira da Silva (Ferreira, 2007) has been computed from the published results graph showing 13 results out of 20 brain scans, and so is not directly comparable with ours. Similarly, results from (Greenspan *et al.*, 2006) were obtained using 18 out of 20 brain scans.

Table 3.2 Average Dice index for various segmentation methods on the newer IBSR 18 normal brain scan, based on published data. FCM+t is our initialization.

Method	WM	GM	CSF
KVL (van Leemput <i>et al.</i> , 1999)	0.857	0.786	0.164
CGMM (Greenspan <i>et al.</i> , 2006)	0.847	0.789	0.219
FCM+t	0.874	0.900	0.296
Local Linear	0.886	0.918	0.344

Also, our initialization procedure labels some pixels on the brain circumference as CSF, and our method has a tendency to expand them.

We have also tested our algorithm, without any pre-processing, on the newer, higher resolution IBSR database with slices 1.5mm thick (Worth, 1996). For this experiment, we set $w = 91$. The average Dice indices that we have obtained are shown in Table 3.2, and they compare favorably with those reported in (van Leemput *et al.*, 1999) and (Greenspan *et al.*, 2006). The general behavior of our method on the newer dataset is similar to that observed on the classic IBSR dataset. Sample results are shown in Fig. 3.10.

3.4 Discussion and conclusions

Our new segmentation algorithm succeeded in outperforming state-of-the-art results on real-world MR brain scan databases. To our knowledge, this is the first time that a level set method has been used for the automatic segmentation of MRI brain tissues with convincing results. Desirable characteristics of our method include: 1) its simplicity, i.e. there is no need to register a sophisticated brain template; 2) its level set implementation, which allows use of the spatial coherence of the voxels to achieve more natural clusters; and 3) its local linear region representatives, which allows moderate-level INU to be taken into account.

During our experiments, we found that our method neither under- nor over-segments the WM, on average. However, we noted that it has a tendency to over-segment the WM around the thalamus and to under-segment it around the pons and the medulla oblongata (e.g. Fig. 3.9). In those regions, the global intensity level of the voxels does not seem to provide enough information for the best segmentation. This limitation suggests that the use of some spatial prior may be required to achieve a segmentation quality that matches those of the experts. However, in order to keep the method generic when using such prior knowledge, care should be taken not to bias the results toward a template. This is a subject for further research.

It has been observed that some of the segmentation results presented in this section have an irregular contour, despite the curvature term in (3.11). This is primarily related to the fact that the parameter ν has been set to a small value. This leads, from our experience, to the results that best match the expert segmentations. The performance of the algorithm with respect to the value of ν is depicted on Fig. 3.3. Visually, the regularity of the contour grow with ν .

All the simulated and real-world scans used in this work were 1.5T, which is representative of today's clinical environment. In contrast, the increased popularity of ultra high-field imaging may results in additional challenges. In particular, an increased magnetic field is associated with more pronounced RF coils-induced inhomogeneity and a stronger bias field. Other authors handle this case explicitly and have presented interesting results on images acquired with field as high as 7T (Li *et al.*, 2008a). On the other hand, with our method if the INU level passes a certain threshold, the global initialization scheme depicted in sec. 3.2.9 may provide initial contours that are too far from the desired segmentation, which would compromise the final results. A localized initialization scheme may be more appropriate. It may also be advisable to use the proposed method in conjunction with some prospective or retrospective correction methods, such as (Sled *et al.*, 1998; Deoni *et al.*, 2006), for the most difficult cases. Nevertheless, the proposed method will likely remain relevant for some time due to the large amount of standard-field MRI scans that have been already recorded and because the 1.5T scanners are commonly used today.

A slice-by-slice implementation has been used throughout this work with satisfactory results. Most of the time, segmentation is consistent from one slice to another. However, because the quality of the segmentation is only linked to the distribution of the voxel intensities on each individual slice, it is possible that the quality of the result will vary greatly from one slice to the other if the dataset has large inter-slice intensity variability e.g. IBSR subject 5_8. This effect can be minimized using several strategies. On the one hand, it may be possible to use a better normalization method during volume acquisition. On the other hand, a fully 3D implementation of the local linear model, which is planned in future work, may also improve slice-to-slice coherence. Because such an implementation will treat the segmentation boundaries as 3D surfaces, it will ensure a certain level of inter-slice spatial coherence. Also, it is expected that this will lead to a faster convergence rate. Finally, a fast narrow band implementation of our method will make it possible to embed it in a real-world application.




Figure 3.4 Sample segmentations obtained with our method on the BrainWeb synthetic scans with a noise level of 3%. *Rows 1 and 2:* slices are 1mm thick and INU is 40%. *Rows 3 and 4:* slices are 3mm thick and INU is 20%. *Columns, from left to right:* input data, reference object used by the simulator, our results, and the difference between our results and the reference. Slices 3mm thick are significantly harder to segment due to increased PVE.

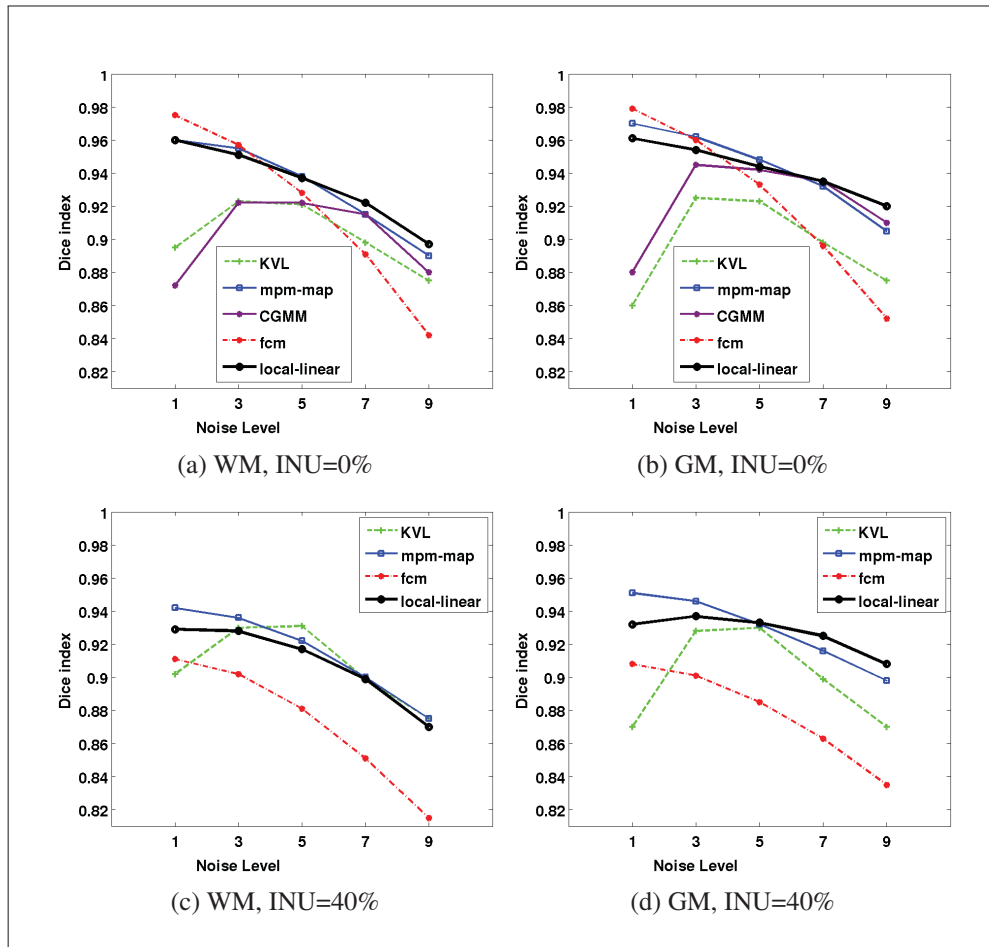


Figure 3.5 Results obtained on the 1mm BrainWeb data for various noise level. Top row: INU=0%, bottom row: INU=40%.

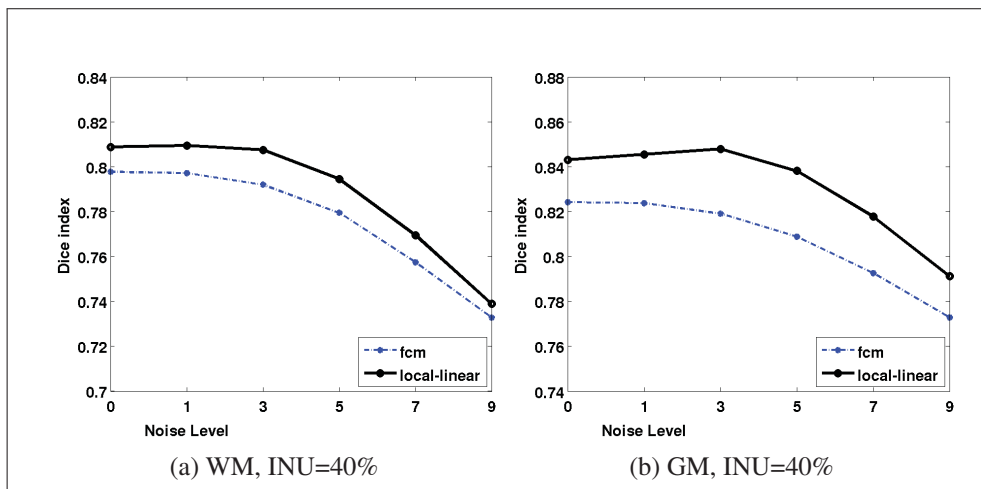


Figure 3.6 Results obtained on the 3mm BrainWeb data, INU=40%.




Figure 3.7 Sample segmentation for the IBSR dataset (Worth, 1996). Left: input data. Middle: reference segmentations. Right: our results.




Figure 3.8 Performance index for the 20 normal brain of the IBSR.
Above: WM, below: GM.




Figure 3.9 Sample result for subject 17_3 of the IBSR database. Top row: input image, reference segmentation, and our result. The two bottom rows illustrate the progression of our local linear model: initialization, after 30 iterations, and our final result. The zoom corresponds to the region selected in the top row.




Figure 3.10 Sample segmentation for the newer IBSR dataset (Worth, 1996). Left: input data. Middle: reference. Right: our results.



CHAPTER 4

3D VESSEL DETECTION FILTER VIA STRUCTURE-BALL ANALYSIS

David Rivest-Hénault¹ and Mohamed Cheriet¹,

¹ Département de génie de la production automatisée, École de Technologie Supérieure,
1100 Notre-Dame Ouest, Montréal, Québec, Canada H3C 1K3

Paper submitted to IEEE Transaction on Image Processing in April 2012

Abstract

Curvilinear structure detection filters are crucial building blocks in many medical image processing applications, where they are used to detect important structures, such as blood vessels, airways, and other similar fibrous tissues. Unfortunately, most of these filters are plagued by an implicit single structure direction assumption, which results in a loss of signal around bifurcations. This peculiarity limits the performance of all subsequent processes, such as understanding an angiography acquisitions, computing an accurate segmentation or tractography, or automatically classifying image voxels. This paper presents a new 3D curvilinear structure detection filter based on the analysis of the *structure ball*, a geometric construction representing second order differences sampled in many directions. The *structure ball* is defined formally, and its computation on a discrete image is discussed. A contrast invariant diffusion index easing voxel analysis and visualization is also introduced, and different *structure ball* shape descriptors are proposed. A new curvilinear structure detection filter is defined based on the shape descriptors that best characterize curvilinear structures. The new filter produces a *vesselness* measure that is robust to the presence of X- and Y-junctions along the structure by going beyond the single direction assumption. At the same time, it stays conceptually simple and deterministic, and allows for an intuitive representation of the structure's principal directions. Sample results are provided for synthetic images and for two medical imaging modalities.

Keywords: Vessel detection, vesselness, vascular system, angiography, CT, CTA, MRI, TOF MRA.

4.1 Introduction

The visualization and study of curvilinear structures is of the utmost importance in medical imaging, where such structures indicate the presence of veins, arteries, bronchi, fibers or other

important functional structures in 3D medical images. Arteriogenesis and stenosis are of particular interest to physicians since they can reveal many severe conditions. Automatic detection of curvilinear structure is also an important component in many medical imaging applications, such as: vessel tracking (Schneider and Sundar, 2010), segmentation (van Bemmelen *et al.*, 2003) and analysis (Frangi *et al.*, 2001), and image-guided surgery (Ruijters *et al.*, 2009; Rivest-Hénault *et al.*, 2012).

Different types of curvilinear structure detection approaches have been proposed in the literature, mostly within the context of blood vessel detection in medical imaging. The most well known is certainly Frangi's method (Frangi *et al.*, 1998, 2001), which belong to the Hessian analysis category (Lorenz *et al.*, 1997; Sato *et al.*, 1998; nero C. and P., 2003; Manniesing *et al.*, 2006b). The assumption is that when a voxel belongs to a salient cylindrical region, the smallest eigenvalue of its Hessian matrix is much smaller than the other two, in term of absolute values. Application of these filters to a 3D image results in the output of a smooth *vesselness* measure for each voxel. These methods are conceptually simple, easy to implement, and reasonably fast, which probably explains their popularity.

Different methods exploit phase congruency models to detect lines, roof edges, or other structures. A recent article by Obara *et al.* (Obara *et al.*, 2012) introduced the phase congruency tensor model for the detection of curvilinear features in 2D images. Their method is based on the analysis of the application of the log-Gabor quadrature filter (Field, 1987) applied in many directions, and achieves brightness- and contrast-invariant detection. Other local curvilinear structure detection operator include those based on the Beamlet transform (Berlemont and Olivo-Marin, 2010), and line integral (Yuan *et al.*, 2011). The flux operator, defined as the integral of the image gradient perpendicular to an arbitrary surface, as been used for the detection of various shapes in medical imaging (Vasilevskiy and Siddiqi, 2002; Law and Chung, 2009). In particular, the spherical flux operator (Bouix *et al.*, 2005; Law and Chung, 2009), and the more specific oriented flux operator (Law and Chung, 2008) are used for vessel segmentation. An interesting aspect of these operators is that they are often used with complementary geometric methods such as the level-set and minimal path methods (Sethian, 1999; Cohen and Kimmel, 1997; Vasilevskiy and Siddiqi, 2002; Benmansour *et al.*, 2009; Law and Chung, 2010).

Explicit modeling of the vessel structure has also been proposed. In this case, a certain shape, usually a cylinder with a circular (Friman *et al.*, 2010), Gaussian (Noordmans and Smeulders, 1998; Wörz and Rohr, 2007), or elliptical (Krissian *et al.*, 2006) cross-section, is fitted to neighboring pixels using an optimization process. Because of the computational cost involved,

these approaches have mostly been limited to tracking, where the number of voxel tested is smaller than in full volume processing. Nevertheless, their performance for interactive vessel segmentation is well demonstrated (Schaap *et al.*, 2009a).

A limitation common to the above approaches is that they make the implicit assumption that a single structure is present locally, which limits the performance at fibers crossing or vessel junction (Lorenz *et al.*, 1997; Frangi *et al.*, 1998; Sato *et al.*, 1998; Law and Chung, 2008; Friman *et al.*, 2010; Obara *et al.*, 2012). An ad hoc criterion is used in (Friman *et al.*, 2010) to detect junctions during tracking, but this approach does not seem applicable for general whole volume vessel detection. The method recently introduced in (Wong *et al.*, 2012) has been designed to handle bifurcations naturally, but so far is limited to interactive vessel lumen segmentation applications. The local line integral approach of Yuan *et al.* (Yuan *et al.*, 2011) does not assume a single dominant direction, and gives good results around bifurcations on 2D images. A probabilistic formulation was proposed by Qian *et al.* in (Qian *et al.*, 2009) to overcome the single structure limitation. Instead of computing image derivatives with a convolution operator, as in the Hessian-based methods, they consider a spherical intensity profile, sampled at many angles and radii. Vessel detection is performed using an information theory criterion. The response of the filter is improved at vessel junctions, but we have found it prone to structural noise in proximity to other structures, such as 3D surfaces. The edge-based approach of Lemaitre *et al.* (Lemaitre *et al.*, 2011) uses a bridging process to complement the missing junction. It is however, limited to 2D images. Machine learning-based methods can also give promising results, but their specialized feature set limits them to very specialized applications, such as coronary artery segmentation (Zheng *et al.*, 2011).

The challenges faced by curvilinear structure detection methods share some similarities with those faced by the diffusion imaging community. In this field, the Gaussian diffusion tensor (Basser and Pierpaoli, 1996) has been found to be limited in the representation of the fiber crossings that occur in brain tissues (Wiegell *et al.*, 2000). Recently, this problem has been partly solved with HARDI and Q-Ball imaging (Tuch, 2004). In these procedures, the diffusivity characteristic of the tissues is probed using as many as 321 gradient directions (Descoteaux *et al.*, 2009). The visualization and analysis of these rich datasets is challenging, but many solutions have been proposed (Frank, 2002; Descoteaux *et al.*, 2007; Aganj *et al.*, 2010), and now Q-Ball and HARDI imaging are becoming more common. We believe that a similar approach is attractive for vessel detection.

In this paper, we define a new vesselness measure based on *structure ball* analysis. The *structure ball* (SBall) is a construction made by evaluating second order numerical difference in

many directions. Since it does not assume any specific distribution, it allows for the representation of structure crossings and junctions. The added complexity makes the analysis more difficult than in Hessian-based vessel detection. Thus, in addition to the introduction and definition of the *structure ball*, the definition of a contrast-invariant structure diffusivity index, and the definition of criteria for vessel detection are the main contributions of this work. The resulting vesselness measurement has good behavior at junctions, and, at the same time, has good discriminative power. In addition, the algorithm can output structural information that can be useful in a potential subsequent visualization, tracking or segmentation step. This methodology is detailed in the next section. Experiments on synthetic and clinical images are presented in section 4.3. Finally, section 5.7 includes a discussion and provides our conclusion.

4.2 Methodology

Detecting curvilinear structures in 3D intensity images, such as CTA with blood pool contrast injection or time-of-flight (TOF) magnetic resonance angiography (MRA) in medical imaging, requires assuming a certain set of hypotheses regarding their nature. The structures of interest depicted in the 3D images are assumed to be mostly curvilinear with roughly circular cross-sections. They might also join or cross each other in certain locations. Although it is not strictly necessary, we restrict the discussion to the cases where at most two vessels are present in one voxel location. This is reasonable in the context of vessel-like structure detection, where X-, T- or Y-shaped bifurcations are much more common than star-like patterns. Since it is assumed that no functional imaging is available, any directional information regarding the underlying tissues must be inferred from the structural image voxels.

The proposed vessel detection filter makes use of a directional second order difference operator, computed in many directions and at every pixel location, to capture the intensity characteristic of an image. Here, the main assumption is that for a bright vessel on a darker background, the second order derivative will have a high absolute value in directions perpendicular to the vessel direction, and a very small value in the parallel direction (Frangi *et al.*, 1998; Sato *et al.*, 1998). As will be discussed in the following sections, the situation is slightly more complex at junctions. In this work, the SBall is defined as the set of all directional differences, as presented in section 4.2.1. In section 4.2.2, we propose to represent this rich directional information using band-limited spherical harmonics (SH). Using SH provides additional regularity, slightly reduces the volume of the data, and makes it possible to compute quantities that are independent of the exact sampling directions. In its raw form, the SBall is difficult to visualize, since the principal direction of an underlying vessel, indicated by a very small second derivative, is ob-

scored by the predominance of the large values, as depicted in Fig. 4.2. The contrast invariant *diffusivity* index proposed in section 4.2.3 enhances visualization, and simplifies our understanding of SBall structures. This *diffusivity* index might also be useful in subsequent tracking or segmentation processes, potentially similar to those in (Descoteaux *et al.*, 2009; Descoteaux and Deriche, 2008; Law and Chung, 2009). Section 4.2.4 presents SBall shape descriptors that are potentially useful for the detection of curvilinear structures. Then, in section 4.2.5, the best descriptors are combined and used to define a smooth vesselness measure that is very discriminative, but also takes into account the potential presence of vessel junctions. Finally, multiscale integration is discussed in section 4.2.6, and section 4.2.7 consist of a discussion on parameter selection.

4.2.1 The *structure ball*: a local structure model

Let $I \subset \mathbb{R}^3$ be the image under consideration. The local structure around the voxel at spatial position \mathbf{x} is inferred by numerically computing the directional second order difference in N different directions. The set of all these directional differences constitutes the SBall. The following simple difference scheme is used:

$$D^2[I, \mathbf{x}, \mathbf{v}] = \frac{I(\mathbf{x} - h\mathbf{v}) - 2I(\mathbf{x}) + I(\mathbf{x} + h\mathbf{v})}{h^2}, \quad (4.1)$$

where $\mathbf{v} = (v_x, v_y, v_z) \in \mathbb{R}^3, \|\mathbf{v}\| = 1$ is a unit vector, and h is the step size. Since the image values at $I(\mathbf{x} \pm h\mathbf{v})$ do not generally correspond to grid points, an image interpolator must be used. Throughout this work, a trilinear interpolator has been used, but other choices are possible (Meijering *et al.*, 1999).

The set of sampled directions constituting the SBall are defined by the vertices of an N -order tessellation of the icosahedron. At each tessellation level, the triangles of the previous level are subdivided into four parts, as depicted in Fig. 4.1. Because the icosahedron and the $D^2[I, \mathbf{x}, \mathbf{v}]$ operator are symmetrical, only the vertices from the north hemisphere are necessary. This results in $\{6, 21, 81, 321\}$ sampling directions for tessellation of order 0 to 3. This choice of sampling directions results in more evenly spaced directional vectors than simply subdividing the spherical coordinate axis, and is the most common choice in Q-Ball imaging (Frank, 2002; Tuch, 2004; Descoteaux *et al.*, 2007). Hessian-based methods infer the image structure direction with 6 directional second order filters (Sato *et al.*, 1998; Frangi *et al.*, 1998, 2001; nero C. and P., 2003; Manniesing *et al.*, 2006b). Thus, in order to provide superior directional accuracy, the rest of the discussion is focused on sampling icosahedrons of order 1 or more.




Figure 4.1 a) Surface defined by a first order subdivision of the icosahedron.
b) Subdivision pattern — Each triangle (solid lines) is divided into four (dashed lines).

For discrete images, the step size h in (4.1) needs to be adjusted as a function of the number of sampling directions. We use two conditions: 1) $h \geq \sqrt{3}$, which ensures that the differentiation step is at least one full voxel in all directions, and 2) the number of sampled voxels must be greater than or equal to twice the number of sampling directions, in order to avoid redundancy. Let us denote by \mathcal{V}_M the set of all sampling directions defined by the M -order tessellation of the icosahedron, and by $\mathcal{N}(\mathcal{V}_M, h)$ the set of all pixels used by the interpolator in all applications of (4.1) for a certain choice of \mathcal{V}_M and step size h . Obviously, the cardinality of that set, $|\mathcal{N}(\mathcal{V}_M, h)|$, depends on the choice of interpolator. For a linear interpolator, we have $\mathcal{N}(\mathcal{V}_M, h) \subseteq \{\mathbf{y} \mid \forall \mathbf{v} \in \mathcal{V}_M \text{ and } \mathbf{y} \in I \text{ if } |\mathbf{y} - (\mathbf{x} + h\mathbf{v})| < 1\}$ for some \mathbf{x} . The second criterion is thus $|\mathcal{N}(\mathcal{V}_M, h)| \geq 2|\mathcal{V}_M|$. Taking into account the two criteria, we found $h = \{\sqrt{3}, 2.62, 5.01\}$ for icosahedron orders $N = \{1, 2, 3\}$ respectively. Thus, there is a trade-off between spatial and angular resolution.

4.2.2 Band-limited spherical harmonics representation

Keeping track of the 21 or more measurements of (4.1) can be tedious and is not memory efficient. In addition, the discrete nature of the SBall makes it difficult to estimate quantities at any arbitrary angle. In this context, spherical harmonics (SH) are used as a smooth, continuous, and compact (for $N \geq 1$) representation of the SBall. SH form the basis for functions on the unit sphere, which can be considered as the spherical analog to the regular Fourier basis. The SH basis are increasingly used in HARDI imaging (Descoteaux *et al.*, 2007), engineering (Garboczi, 2002), and in computer graphic (Sloan, 2008). The SH functions, denoted by Y_l^m , are defined as follows (Descoteaux *et al.*, 2007):

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{im\phi}, \quad (4.2)$$

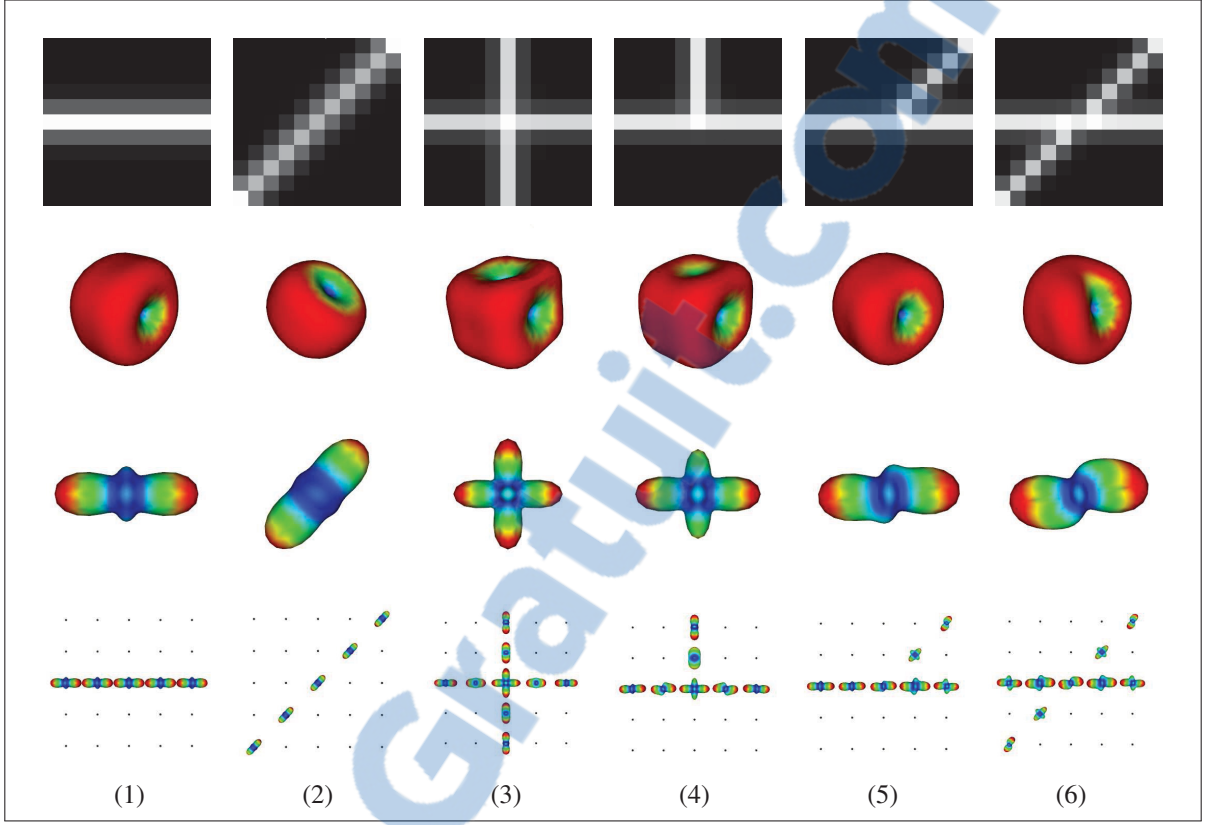


Figure 4.2 Illustration of the shape of the SBall and DBall for various configurations. *Top row:* Center slices of different low-resolution $11 \times 11 \times 7$ volumes. *Second row:* SBall representation of the center voxels — see section 4.2.1. *Third row:* Corresponding DBalls — see section 4.2.3. *Bottom row:* DBall grids. The DBalls are scaled by their vesselness value. Small spheres indicate voxels with vesselness < 0.1 .

where $\theta \in [0, \pi]$ is the elevation angle, $\phi \in [0, 2\pi]$ is the azimuth and P_l^m is an associated Legendre polynomial (see, for example, (Garboczi, 2002; Aganj *et al.*, 2010) for a listing). As the SBall structure is real and symmetrical, a modified basis with those properties can be used (Descoteaux *et al.*, 2007):

$$Y_j = \begin{cases} \sqrt{2} \cdot \text{Re}(Y_k^m) & \text{if } -k \leq m < 0 \\ Y_k^0 & \text{if } m = 0 \\ \sqrt{2} \cdot \text{Im}(Y_k^m) & \text{if } 0 < m \leq k \end{cases}, \quad (4.3)$$

with the index j defined as $j := j(k, m) = (k^2 + k + 2)/2 + m$ for $k = 0, 2, 4, \dots, l$, and $m = -k, \dots, 0, \dots, k$. A function defined on the unit sphere can then be approximated by:

$$S(\theta, \phi) = \sum_{j=1}^R c_j Y_j(\theta, \phi) \quad (4.4)$$

where $R = (l+1)(l+2)/2$ depends on the order l of the SH basis chosen to represent the signal. Thus $R = 1, 6, 15, 28, \dots$ for $l = 0, 2, 4, 6, \dots$. There are various means of estimating the coefficients c_j defining the approximation (Frank, 2002; Descoteaux *et al.*, 2007). In this paper, the method of Descoteaux *et al.* (Descoteaux *et al.*, 2007) has been used, and is summarized here for the sake of completeness. Let \mathbf{S} be the $N \times 1$ vector representing the N sampled values of $D^2[I, \mathbf{x}, \mathbf{v}]$ at an image location \mathbf{x} , and \mathbf{C} be the $R \times 1$ vector of the c_j coefficients. Building the matrix \mathbf{B} as follows:

$$\mathbf{B} = \begin{pmatrix} Y_1(\theta_1, \phi_1) & Y_2(\theta_1, \phi_1) & \dots & Y_R(\theta_1, \phi_1) \\ \vdots & \vdots & \ddots & \vdots \\ Y_1(\theta_N, \phi_N) & Y_2(\theta_N, \phi_N) & \dots & Y_R(\theta_N, \phi_N) \end{pmatrix},$$

with (θ_i, ϕ_i) corresponding to the spherical coordinate defined by the i^{th} vertex of \mathcal{V}_M , we have the linear system $\mathbf{S} = \mathbf{B}\mathbf{C}$. This linear system is either well constrained or over constrained when $N > R$, and so the order M of the sampling icosahedron must be chosen considering the order l of the SH representation (since R is a function of l). The coefficients are computed using least squares fitting (Descoteaux *et al.*, 2007) :

$$\mathbf{C} = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \mathbf{S}, \quad (4.5)$$

In HARDI, an SH basis of order 8 or more is commonly used (Descoteaux *et al.*, 2007). Whereas in this case the brain fibers can be arranged in very complex patterns (Wiegell *et al.*, 2000; Frank, 2002), the vessel structures to be detected by the method proposed in this paper are usually simpler. An SH basis of order 4 has been found to adequately represent a signal that presents up to 3 principal orientations, and is assumed to be appropriate for our curvilinear structure detection method. Since an SH of order 4 has 15 coefficients, the estimation (4.5) requires the use of a sampling icosahedron of order 1 or more so that the problem of estimating the SH coefficients is overconstrained.

4.2.3 Contrast-invariant diffusivity index

In an ideal context, a voxel traversed by a single vessel is characterized by an SBall with a high amplitude everywhere, except in one direction where its amplitude is small, or even zero. This intuitive description has the disadvantage of being contrast-dependent, in the sense that it is sensitive to a rescaling of the image intensity, which complicates the definition of a robust descriptor. In addition, the raw SBalls are difficult to visualize, since the most probable structural directions are hidden behind the large amplitude regions that characterize the vessel's cross-sectional directions. So, we define a contrast-invariant diffusivity index, which is maximal in the most probable direction, as follows:

$$\mathcal{D}(x) = \exp\left(-\left(\frac{x}{S^*}\right)^2\right), \quad (4.6)$$

where $\gamma \in \mathbb{R}$, and $S^* = \text{average}[S(\theta, \phi)]$. For computational efficiency, the approximation $S^* \approx \frac{1}{N} \sum_i [D(\theta_i, \phi_i)]$ has been used throughout this paper. Also, we slightly abuse the notation, and use $\mathcal{D}(\theta, \phi) \equiv \mathcal{D}[S(\theta, \phi)]$. The non linear mapping introduced in (4.6) also helps to limit the influence of nearby outliers structures, such as bones, that might result in very high second derivative values. Indeed, $\mathcal{D}(\theta, \phi) \approx 0$ when $S(\theta, \phi) \geq 3S^*$. In the sections below, we denote by DBall the application of (4.6) on all the directions of an SBall. This mapping is demonstrated in Fig. 4.2.

4.2.4 Geometric descriptors

In this section, a set of geometric descriptors that can be used to discriminate curvilinear structure voxels from background voxels is proposed. The definitions presented here are related to concepts previously introduced in (Frangi *et al.*, 1998; Sato *et al.*, 1998; Bouix *et al.*, 2005; Law and Chung, 2008) and also to those in (Basser and Pierpaoli, 1996; Frank, 2002; Tuch, 2004). However, they have been adapted to the SBall structure, and the single direction assumption has been dropped. Instead, the key hypothesis in this section is that, when a voxel is traversed by one or two vessels, its SBall presents a related number of strong minima. Furthermore, it is assumed that the main structural directions can be well represented on a plane cutting the SBall at a certain angle. This assumption is reasonable when the curvature of the underlying structure is not very high compared to its radius, which is usually the case, for example, with blood vessel in medical imaging.

4.2.4.1 Ratio descriptors

For curvilinear structures, it is clear that value of the SBall is large, on average, with some regions of low value, indicating the presence of vessels (see Fig. 4.2). So, the ratio of minimum to average values of the SBall is a good indicator of the presence of one or two vessels in the surrounding area. Another possibility would be to consider the ratio of the minimum to maximum values. In addition, since the evaluation of the minimal SBall value is likely to be noisy when the structural direction is not perfectly aligned with the sampling direction, we also consider squaring the ratios. This reduces the sensitivity for small ratio values, and increase it for large ones. Let

$$\begin{aligned}\bar{D} &= \text{average}_i(|D(\theta_i, \phi_i)|) \\ D^+ &= \max_i(|D(\theta_i, \phi_i)|) \\ D^- &= \min_i(|D(\theta_i, \phi_i)|)\end{aligned}\tag{4.7}$$

be the average SBall value, the minimal value, and the maximal value respectively. We define our ratio descriptors as:

$$\frac{D^-}{\bar{D}}, \quad \frac{D^-}{D^+}, \quad \left(\frac{D^-}{\bar{D}}\right)^2, \text{ and } \left(\frac{D^-}{D^+}\right)^2.\tag{4.8}$$

Since $D^- \leq \bar{D} \leq D^+$, the ratios are always in the $[0, 1]$ range. These descriptors are indications of how close the structure is to a ball. However, they do not accurately discern a curvilinear structure from a disk-shaped ball.

4.2.4.2 Oriented bounding box

From Fig. 4.2, it is clear that the aspect of the oriented bounding box (OBB) of a DBall offers a good indication of the presence of one or two vessels in the surrounding area. Let \mathbf{v}_1 , \mathbf{v}_2 , and \mathbf{v}_3 be the three dimensions of the OBB, such that $|\mathbf{v}_1| \geq |\mathbf{v}_2| \geq |\mathbf{v}_3| \geq 0$ (see Fig. 4.3). If one vessel is present, $|\mathbf{v}_1| \gg |\mathbf{v}_2| \geq |\mathbf{v}_3|$ is observed. If there are two vessels, only one direction is small. So, we define our bounding box aspect descriptor function as follows:

$$\text{OBB} = \frac{\mathbf{v}_3}{\mathbf{v}_1}.\tag{4.9}$$

4.2.4.3 Fractional anisotropy

Generalized fractional anisotropy (GFA) was proposed by Tuch in (Tuch, 2004) as an extension of Basser and Pierpaoli's fractional anisotropy (Basser and Pierpaoli, 1996), and was used to ease the visualization of the Q-Balls. This measure is as follows:

$$\text{GFA}(\mathbf{x}) = \frac{\sqrt{\sigma^2(\mathbf{x})}}{\text{RMS}(\mathbf{x})}, \quad (4.10)$$

where $\sigma^2(\mathbf{x})$ and $\text{RMS}(\mathbf{x})$ represent the variance and root mean square value of the ball centered at \mathbf{x} . The main idea is that a strongly anisotropic structure has a large variance, and a constant function has no variance. Similar concepts were used by Frank for defining his fractional multi-fiber index (Frank, 2002). Because vessels are strongly anisotropic, GFA can be used as a vessel indicator on the SBall or the DBall. For this reason, we propose to compute GFA on both structures for curvilinear structure detection. We denote by GFA_S GFA computed on the SBall values, and by GFA_D GFA computed using the DBall values.

Nonetheless, GFA can also be high for other anisotropic structures. For this reason, we propose a more discriminative planar fractional anisotropy (PFA) measure, computed along the largest circle of the DBall. This new descriptor is defined as follows:

$$\text{PFA} = \frac{\sqrt{\int_{\mathbf{w} \perp \mathbf{v}_3} (\mathcal{D}(\mathbf{w}) - \bar{\mathcal{D}})^2 d\mathbf{w}}}{\int_{\mathbf{w} \perp \mathbf{v}_3} (\mathcal{D}(\mathbf{w}))^2 d\mathbf{w}}, \quad (4.11)$$

where the integral is defined on the circle that is perpendicular to the smallest oriented bounding box direction, \mathbf{v}_3 , as depicted in Fig. 4.3. Here, $\bar{\mathcal{D}}$ denotes the mean value along the same circle. As this formulation is defined continuously, the values between the sampling icosahedron directions must be interpolated. This can be done using the SH representation introduced in section 4.2.2. Numerically, the PFA descriptor is computed by sampling along the great circle defined by \mathbf{v}_3 . It can be shown that both GFA and the PFA are always in the $[0, 1]$ range.

4.2.4.4 Flux

The descriptors introduced up to now have two limitations for vessel detection: 1) they cannot discriminate between a bright structure on a dark background, or the converse, and 2) they give no indication of the saliency of the structure. Flux-based descriptors can be used for this purpose.




Figure 4.3 Oriented bounding box of a DBall (*left*), and computation of the $\mathcal{A}(\mathbf{x})$ function on the circle defined by the \mathbf{v}_3 vector (*right*).

In a vector field, the spherical flux is defined as the integral of the vector quantity that flows out of a spherical surface, in the directions perpendicular to this surface. This quantity has been used in many imaging applications, including vessel detection (Bouix *et al.*, 2005) and segmentation (Vasilevskiy and Siddiqi, 2002; Law and Chung, 2009). Considering the structure of the SBall induced by (4.1) in section 4.2.1, the spherical flux is well approximated, for a sphere of radius of $r = h/2$, by:

$$\text{SFlux}(\mathbf{x}) \approx \frac{1}{2} \int \int S(\theta, \phi) \cos(\theta) d\theta d\phi. \quad (4.12)$$

Normalizing this flux by the sphere surface $4\pi r^2$, and recognizing that the mean of a function is given, up to a scale factor, by the first SH term in (4.4), we found the following:

$$\text{NSFlux}(\mathbf{x}) = \frac{1}{4\pi r^2} \text{SFlux}(\mathbf{x}) = \frac{1}{4\sqrt{\pi}} c_1, \quad (4.13)$$

where c_1 is the first coefficient in (4.4). The scale factor $\sqrt{4\pi}$ derives from the normalization used in (4.2). This means that, if we have the SH decomposition of the SBall, then the normalized spherical flux at this position has already been computed. The function NSFlux is negative for a bright vessel on a darker background, and positive otherwise. In addition, its modulus is an indication of the structure's saliency.

The optimally oriented flux (OOF) was introduced by Law and Chung (Law and Chung, 2008) as an enhancement to the spherical flux. In their formulation, they consider only the flux passing through the great circle perpendicular to a direction defined by a vector \mathbf{p} , as opposed to the whole sphere surface, as in (4.12). The vector \mathbf{p}^* optimizing the oriented flux is found analytically, and the associated flux quantity is used for curvilinear structure detection.

The oriented flux can easily be computed for any direction using the SBall SH representation, by quickly computing its Funk–Radon transform (FRT). For a spherical point $\mathbf{p} = (\theta, \phi)$, the FRT is the great circle integral of the signal on the circle defined by the plane passing through the origin with normal vector \mathbf{p} . In the SBall context, the FRT is thus equivalent to the oriented flux, up to one multiplicative constant. Descoteaux *et al.* show that the FRT $\mathcal{G}[\mathcal{S}]$ of a function represented using SH basis is (Descoteaux *et al.*, 2007):

$$\mathcal{G}[\mathcal{S}](\mathbf{p}) = \sum_{j=1}^R 2\pi P_{l_j} c_j Y_j(\mathbf{p}).$$

So, it is possible to compute the oriented flux in any direction by a simple summation. Unfortunately, finding the maximum on an SH basis analytically is actually not possible (Aganj *et al.*, 2010), and doing so numerically would cause an additional computational burden. Nevertheless, we have included the following version of a normalized oriented flux descriptor for comparison purposes:

$$\text{NOFlux}(\mathbf{x}, \mathbf{q}) = \frac{1}{4\pi} \sum_{j=1}^R 2\pi P_{l_j} c_j Y_j(\mathbf{q}), \quad (4.14)$$

where \mathbf{q} is the normalized spherical coordinate direction corresponding to \mathbf{v}_3 (see section 4.2.4.2). The derivation of the normalization factor $(1/4\pi)$ is similar to that in (4.13).

4.2.5 Vesselness measure

The proposed geometric descriptors are compared using a set of 14 different patterns representing curvilinear structures (6 patterns from Fig. 4.2) and other interfering shapes (8 patterns from Fig. 4.4). The computed values are presented in Table 4.1 for the curvilinear structure configurations (including bifurcations) and in Table 4.2 for the interfering shapes. As we knew intuitively, all the proposed contrast-invariant indicators are suitable for discerning the near spherical configuration (Config. 11) from the six curvilinear configurations presented. However, it appears that the $(D^-/\bar{D})^2$ ratio was the most sensitive with respect to the other patterns associated with the spherical figure (Configs. 12, 13, and 14). The two disk-like patterns (Configs. 7 and 8) associated with planar surfaces can be distinguished from the curvilinear structure patterns using the PFA. In this case, the PFA is the only descriptor that gives a response that is significantly different for the planar structures than for the curvilinear structures. We found the GFA_S , GFA_D , and OBB descriptors less attractive because they do not produce significantly different output for the interfering patterns. The computed flux (NSFlux and NOFlux) was generally lower for the interfering pattern than for the curvilinear structures, but, because the




Figure 4.4 DBall representation for some interfering patterns. *Top row*: Volume rendering of the test images. *Other rows*: DBall representation for the marked voxel positions. The numbering continues from Fig. 4.2.

flux is mostly representative of the local contrast, it is difficult to determine the significance of this observation. In our setting, we did not find a significant difference between the NSFlux and NOFlux descriptors. We could not identify any shape descriptors that had a clear discriminative power against sharp corners and sharp edge patterns (Configs. 9 and 10), although the PFA is reduced in the latter case. We argue that this is not too problematic in the context of medical imaging, since these configurations are not found in a large proportion of the voxels in biological images.

Based on these observations, we chose three complementary shape descriptors for our new vesselness measure. The ball descriptor $\mathcal{B}(\mathbf{x}) = (D^-/\bar{D})^2$ makes it possible to discriminate ball-like structures, the planar anisotropy descriptor $\mathcal{A}(\mathbf{x}) = \text{PFA}(\mathbf{x})$ enables detection of disk shapes, and, finally, the flux operator $\mathcal{F}(\mathbf{x}) = \text{NSFlux}(\mathbf{x})$ is used to cancel very low contrast structures. Our vesselness measure for detecting bright structures against a darker background

Table 4.1 Values of the proposed shape descriptors computed on the six vessel voxel configurations presented in Fig. 4.2.

Config.	SBall descriptors					DBall descriptors			Flux	
	$\frac{D^-}{D}$	$\frac{D^-}{D^+}$	$\left(\frac{D^-}{D}\right)^2$	$\left(\frac{D^-}{D^+}\right)^2$	GFA _S	OBB	GFA _D	PFA	NSFlux	NOFlux
1	0.000	0.000	0.000	0.000	0.250	0.293	0.397	0.419	-12.752	-14.145
2	0.295	0.251	0.087	0.063	0.216	0.306	0.372	0.380	-10.328	-11.645
3	0.199	0.155	0.040	0.024	0.282	0.219	0.428	0.265	-12.798	-12.228
4	0.102	0.084	0.010	0.007	0.241	0.253	0.380	0.289	-12.798	-13.223
5	0.000	0.000	0.000	0.000	0.268	0.278	0.419	0.388	-12.334	-14.056
6	0.000	0.000	0.000	0.000	0.308	0.262	0.470	0.391	-11.971	-14.022
Min.	0.000	0.000	0.000	0.000	0.216	0.219	0.372	0.265	-12.798	-14.145
Max.	0.295	0.251	0.087	0.063	0.308	0.306	0.470	0.419	-10.328	-11.645

Table 4.2 Values of the proposed shape descriptors computed on the eight non vessel voxel configurations presented in Fig. 4.4. Values in **bold** are significantly different from the vessel voxel configuration in Table 4.1. The criteria for significant difference is either twice the maximum value in Table 4.1 or half the minimum value.

Config.	SBall descriptors					DBall descriptors			Flux	
	$\frac{D^-}{D}$	$\frac{D^-}{D^+}$	$\left(\frac{D^-}{D}\right)^2$	$\left(\frac{D^-}{D^+}\right)^2$	GFA _S	OBB	GFA _D	PFA	NSFlux	NOFlux
7	0.000	0.000	0.000	0.000	0.503	0.188	0.597	0.041	-11.477	-13.752
8	0.000	0.000	0.000	0.000	0.503	0.188	0.597	0.041	-4.865	-5.829
9	0.225	0.139	0.050	0.019	0.461	0.530	0.592	0.508	-9.639	-13.811
10	0.000	0.000	0.000	0.000	0.479	0.341	0.519	0.149	-7.934	-9.088
11	0.942	0.915	0.888	0.838	0.035	1.000	0.048	0.048	-8.782	-8.599
12	0.451	0.329	0.203	0.108	0.307	0.428	0.476	0.442	-4.452	-5.625
13	0.560	0.462	0.314	0.213	0.183	0.380	0.315	0.305	-5.799	-6.747
14	0.735	0.675	0.540	0.456	0.092	0.580	0.161	0.170	-7.250	-7.694

combines the three measures previously defined in a smooth function, as follows:

$$\mathcal{V}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathcal{F}(\mathbf{x}) > 0 \\ V(\mathbf{x}) & \text{otherwise} \end{cases}, \quad (4.15)$$

with

$$V(\mathbf{x}) = \exp\left(\frac{-(\mathcal{B}(\mathbf{x}))^2}{2\beta^2}\right) \cdot \left(1 - \exp\left(-\frac{(\mathcal{A}(\mathbf{x}))^2}{2\alpha^2}\right)\right) \cdot \left(1 - \exp\left(-\frac{(\mathcal{F}(\mathbf{x}))^2}{2\gamma^2}\right)\right). \quad (4.16)$$

The vesselness is zero if the voxel is, on average, darker than the surrounding area, as indicated by a positive spherical flux. Obviously, this condition can be reversed to detect a dark structure

on a brighter background. The vesselness measure is close to zero if the structure appears like a ball ($\mathcal{B}(\mathbf{x}) \approx 1$), or a disk ($\mathcal{A}(\mathbf{x}) \approx 0$), or has very low contrast ($\mathcal{F}(\mathbf{x}) \approx 0$), and is close to unity otherwise.

4.2.6 Multiscale integration

The size of the structures that can be captured by an SBall is limited by the step size of the difference operator h in (4.1). We envision two possibilities for detecting the structures of variable scales: 1) scaling the sampling icosahedron and the step size h , and 2) embedding the SBall analysis in a scale-space framework. Since the former approach appears less practical, as the computational time required by an accurate implementation will increase with the square of the scale of the structure, we develop the latter idea here, using the scale-space formalism of Lindeberg (Lindeberg, 1996). Within this framework, structures of scale σ are detected by using the following scale-normalized second order differences instead of the one defined in (4.1):

$$D_\sigma^2[I, \mathbf{x}, \mathbf{v}] = \sigma^2 D^2[I_\sigma, \mathbf{x}, \mathbf{v}] \quad , \quad (4.17)$$

where $I_\sigma = I * G(\sigma)$ is the input image convolved with a Gaussian filter $G(\sigma)$ of σ standard deviation. With this formulation, the only other quantity that needs to be renormalized is the net flux, as computed in (4.13). Flux is a first order quantity, and requires a different scale normalization than a second order quantity, as in (4.17) (Lindeberg, 1996):

$$\mathcal{F}_\sigma(\mathbf{x}) = \frac{1}{\sigma} \text{NSFlux}(\mathbf{x}) \quad (4.18)$$

By assuming a discrete range of scales σ_i , and by denoting by $\mathcal{V}_{\sigma_j}(\mathbf{x})$ the vesselness response for some scale σ_j , we define the multiscale vesselness response as

$$\mathcal{V}_{\text{multi}}(\mathbf{x}) = \max_i \mathcal{V}_{\sigma_i}(\mathbf{x}) \quad . \quad (4.19)$$

The response of $\mathcal{V}_{\sigma_j}(\mathbf{x})$ and $\mathcal{V}_{\text{multi}}(\mathbf{x})$ across a range of scale is shown in Fig. 4.5.

4.2.7 Parameter selection

The two first terms of $V(\mathbf{x})$ in (4.16) are independent of the overall image contrast, and so the parameters β and α can be fixed with only geometrical considerations. The last term does depend on the image contrast. This dependency can be mitigated by using a value $\gamma = \xi \cdot \mathcal{F}^*$,

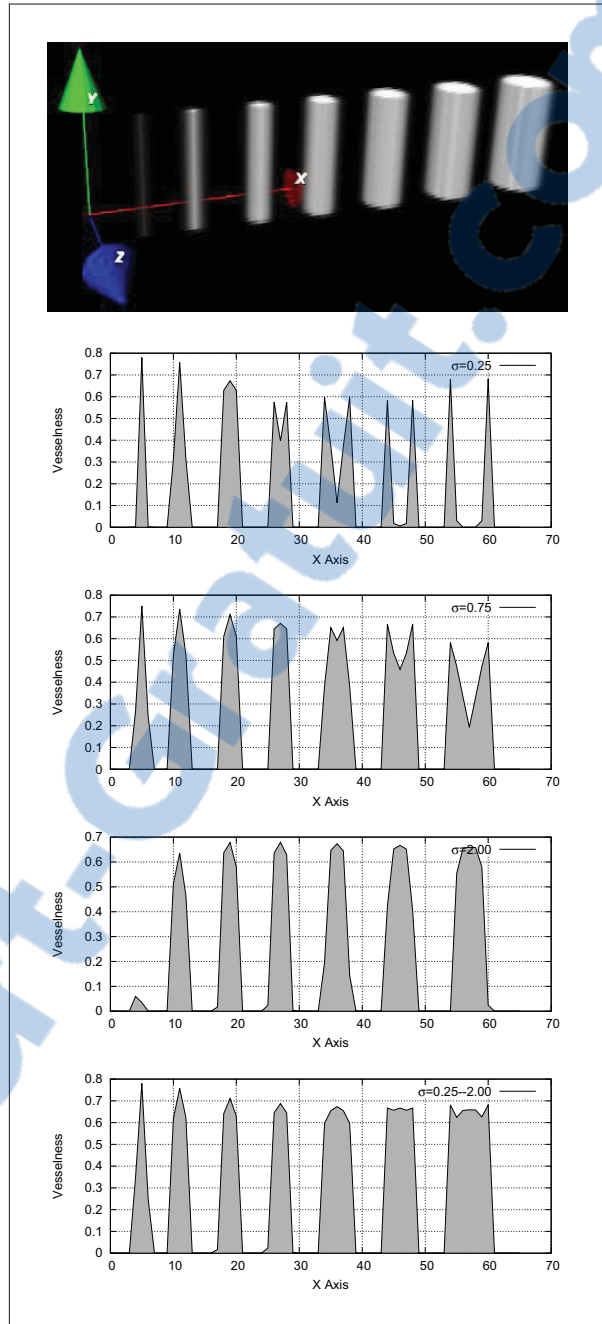


Figure 4.5 Multiscale response for simple cylindrical patterns with radii of $\{0.5, 1.0, \dots, 3.5\}$ voxels in a $16 \times 16 \times 66$ volume. *From top to bottom:* Synthetic input volume, intensity of the vesselness signal along the line defined by $\{y = 8, z = 8\}$ for scales $\sigma = \{0.25, 0.75, 2.00\}$, and maximal intensity along the same line across the scales $\sigma = \{0.25, 0.50, 0.75, 1.00, 1.50, 2.00\}$.

which is a multiple of the maximal absolute value $\mathcal{F}^* = \max |\mathcal{F}(\mathbf{x})|$ for the whole image.

In an interactive framework, another strategy might be to fix this value based on the contrast observed in certain used-specified regions.

In this paper, parameters β , α , and γ were set by assuming a loss of signal ε for the curvilinear structure configurations presented in Table 4.1 with the higher \mathcal{B} or the lower \mathcal{A} or \mathcal{F} signals ($\mathcal{B}_{\max} = 0.098$, $\mathcal{A}_{\min} = 0.265$ and $\mathcal{F}_{\min} = -10.328$). Also, $\mathcal{F}^* = -12.798$, and so we searched β , α , and ξ , such that:

$$\begin{aligned}\varepsilon &= \exp\left(\frac{-\mathcal{B}^2}{2\beta^2}\right) \\ \varepsilon &= 1 - \exp\left(-\frac{\mathcal{A}^2}{2\alpha^2}\right) \\ \varepsilon &= 1 - \exp\left(-\frac{\mathcal{F}^2}{2(\xi \cdot \mathcal{F}^*)^2}\right),\end{aligned}$$

where \mathcal{F}^* is the maximum spherical flux value in Table 4.1. Reasonable values for ε are in the $[\sqrt[3]{0.50}, 1.00]$ range. Using $\varepsilon = 0.95$, we found $\beta = 0.272$, $\alpha = 0.108$ and $\xi = 0.330$. A ε value closer to 1.0 will provide better sensitivity at the expense of discrimination, and vice-versa for a lower ε value.

4.3 Experimental results

The proposed SBall vessel detection filter has been implemented in C++ as a multithreaded ITK filter. A version of this code will be made freely available for research purposes following publication.

Parameters β , α , and ξ were kept constant for all the experiments, which is an indication of the robustness of the method. In addition, an order 1 sampling icosahedron has been used in all cases. Because of its larger difference step h , which induces some smoothing, the scale σ for the proposed filter in the multiscale configuration is not directly comparable to the scale of Frangi (Frangi *et al.*, 1998) or Sato (Sato *et al.*, 1998) filters. Since, for an order 1 sampling icosahedron $h = \sqrt{3} \approx 2$, we used the equivalence $s = 2\sigma$, where s is Frangi's or Sato's operating scale.

4.3.1 Experiments on synthetic images

Two synthetic images were used to assess the performance of the proposed method. The volumes in this section were generated by drawing the structures in an initial volume with $10\times$ as the indicated resolution, and then scaling it down to the desired size. Consequently, these

images are realistic with respect to the partial volume effect observed in medical imaging. The results are also compared with other established methods, those of Frangi (Frangi *et al.*, 1998), Sato (Sato *et al.*, 1998), and Qian (Qian *et al.*, 2009). In all cases, the parameters have the values recommended by the authors in the original publications. We also compare our method with a flux operator, which is the last term in (4.16), and including the thresholding in (4.15). Finally, the images in this sections were processed at scales $\sigma_i \in \{0.25, 0.50, 0.75, 1.00\}$.

The first test volume represents a $98 \times 32 \times 12$, $I(\mathbf{x}) \in [0.0, 1.0]$, fish bone like structure, that as been designed to test the performance of the method in the presence of vessels joining at various angles ($\theta \in [22.5, 90]\text{deg}$). This test volume is depicted in Fig. 4.6 along with its DBall representation for some key regions. The scalar response of each filter is depicted in Fig. 4.7.

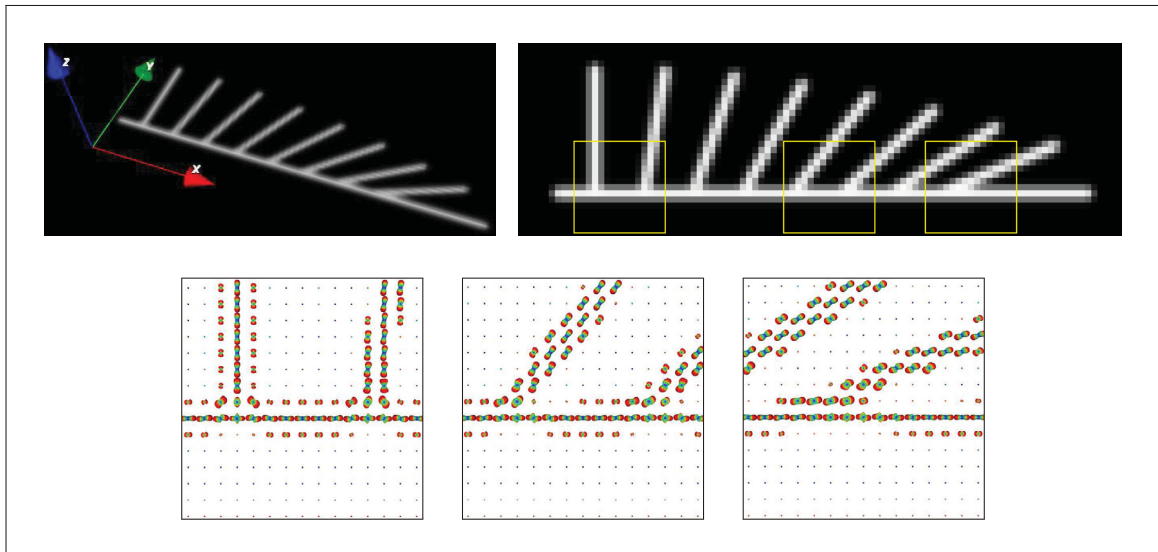


Figure 4.6 Multiscales response for simple 3D vessel pattern at various angles. *Top row:* Synthetic input volume and median slice of the structure. *Bottom row:* Details of the DBalls for the three highlighted areas.

In all cases, the intensity has been normalized so that the maximum vesselness computed by a filter equals 1.0.

The second volume, of size $98 \times 32 \times 12$ voxels and intensity $I(\mathbf{x}) \in [0.0, 1.0]$, has been designed to test the specificity of the vesselness filter. It includes a block, a thin plate, spheres with radii of $r \in [0.75, 3.00]$ voxels, and a cylinder with a radius of 1 voxel. These *structural noise* patterns are meant to represent the various interfering structures that might be present in medical images, such as bones or membranes. The test volume, and the vesselness response on the median slice for the five tested filters are presented in Fig. 4.8.




Figure 4.7 Vesselness filter responses for the slice in Fig. 4.6 for 5 different filters. *From top to bottom: SBall, Frangi's filter, Sato's filter, Qian's filter, and the Flux filter.*

From what can be seen from the fish bone pattern in Fig. 4.7, the proposed vesselness filter is better at preserving the connectivity at vessel junctions than either Frangi's or Sato's filter. However, Qian's filter and the Flux filter show no sign of signal drop at junctions, and there is even a slight signal increase for Qian's filter. Nevertheless, the results presented in Fig. 4.8 suggest that Qian's filter, the Flux filters, and also Sato's filter are much less specific than the two others: they detect the interfering patterns with an intensity that is comparable to that of cylinder detection. The proposed method and Frangi's filter give better results on this volume, and our SBall filter seems to be the best at specificity. Overall, it appears that there is some trade-off between the ability to detect vessels at junctions and filter specificity, and we found that our proposed filter offers the best compromise.

In order to quantitatively assess the performance of each filter, True Positive Rate (TPR) and False Positive Rate (FPR) curves have been plotted. Assuming a known true segmentation




Figure 4.8 Vesselness filter responses for the median slice of the structural noise test volume. *Left to right and top to bottom*: Test volume, and the vesselness for the SBall, Frangi's filter, Sato's filter, Qian's filter, and Flux filter.

$T(x) \in \{0, 1\}$ and some classification threshold t , the TPR and FPR are given by:

$$\text{TPR} = \frac{(\mathcal{V}(x) \geq t) \wedge T(x)}{\#\{T(x) == 1\}}, \quad \text{FPR} = \frac{(\mathcal{V}(x) < t) \wedge T(x)}{\#\{T(x) == 0\}},$$

where $\#$ denotes the cardinality of a set. The curves are then obtained by computing the TPR and FPR for many values in the $t \in [0.0, 1.0]$ range, in terms of normalized vesselness. Intuitively, the TPR can be regarded as the hit rate, that is, the number of voxels that correctly identify the target structure. The FPR can be interpreted as the false alarm rate, or the number of times a filter signals an unwanted structure. For these simple synthetic images, the true segmentations have been assumed to correspond to $I(\mathbf{x}) > 0.5$.

The quantitative results are found in Fig.4.9 for the fish bone pattern, and in Fig. 4.10 for the structural noise patterns. The four graphs presented here tend to corroborate the observations made from the vessel filter intensity images. In fact, with a level of $\varepsilon = 0.95$, we found that the proposed filter is better than the others at almost any given classification threshold for the two test volumes. So, on average, its hit rate of our filter is very high, being only second to the Flux operator, while its false alarm rate is lower than the others, although we acknowledge that Frangi's filter false alarm rate is also low.

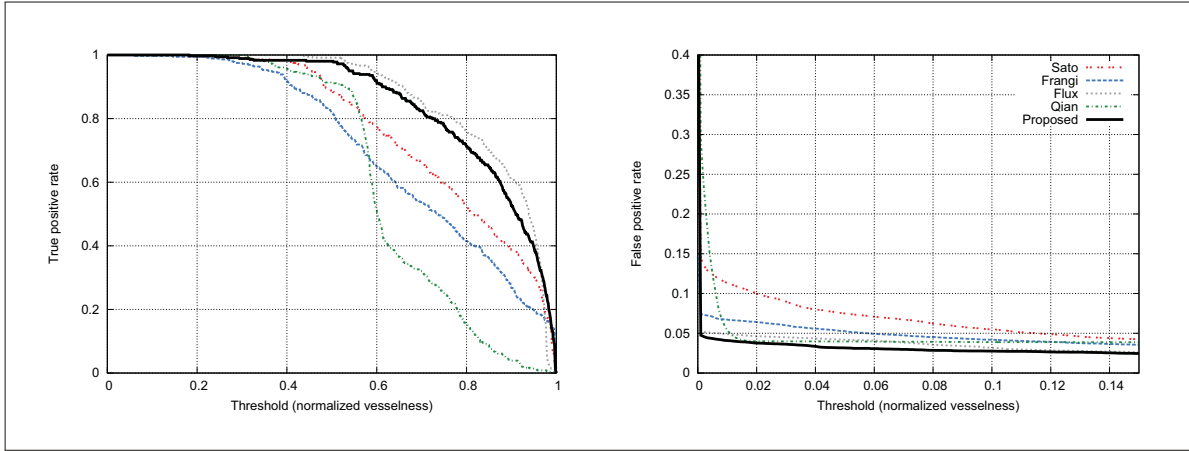


Figure 4.9 True Positive Rate (*left*), and False Positive Rate (*right*) for the fish bone image as a function of the classification threshold.

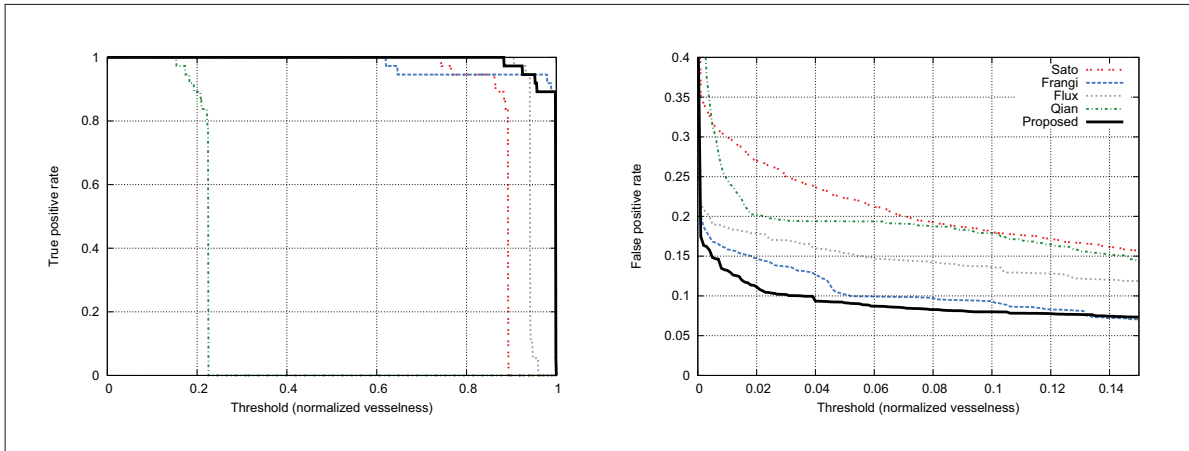


Figure 4.10 True Positive Rate (*left*), and False Positive Rate (*right*) for the structural noise image as a function of the classification threshold.

4.3.2 Experiments on clinical images

The proposed vessel detection filter has been tested on two vastly different clinical datasets: one time-of-flight magnetic resonance angiography (TOF-MRA) scan, and one cardiac computerized tomography angiography (CTA) scan. For comparison, we also included results obtained with Frangi's and Sato's filters.

A $86 \times 152 \times 74$ region of a TOF-MRA scan¹, selected to highlight the filters' differences, has been processed at scales $\sigma_i \in \{0.25, 0.50, 0.75, 1.00\}$. The results for the three vessel detection

¹The TOF-MRA scan tested is distributed with the MRicroGL visualization software, both of which are freely available on the Internet.

filters are shown in Fig. 4.11. The performance of all three filters was relatively good on this dataset. There are few notables differences, however, as indicated by the arrows in Fig. 4.11: **A)** The connectivity at this vessel junction is better preserved by the proposed filter than by Frangi's. Sato's filter has a slightly better response, but it merged all three vessels into one blob. **B)** With the proposed filter, the two vessels are discernible, whereas with the other two filters they are fused. **C)** Sato's filter has the best response for very small vessels, but it also presents the highest signal for unwanted structures **(D)**. For small vessels, the SBall filter, is however significantly better than Frangi's filter.

The performance of the vessel detection filters was also assessed on a more challenging CTA scan. This dataset was made available by Dr. Nagib Dahdah of Ste-Justine Hospital in Montreal. It was acquired on a sane subject after enlightened consent has been given, and after review by the hospital's ethics board. Compared to the other clinical dataset, this scan has less contrast, is more noisy, and presents a larger variety of anatomical structures. Because of the greater image noise, we processed the image at higher scales. Consequently, the same set of scales was used for the proposed SBall filter, for Frangi's and for Sato's filter ($s_i = \sigma_i \in \{0.75, 1.00, 1.50, 2.00\}$). To make the coronary arteries easier to see, the non-heart structures that are far away from the myocardium have been manually masked. The results are presented in Fig. 4.12.

We also found that the proposed filter preserved the junctions of this dataset better than Frangi's filter **(A)**, while having good ability to segment smaller vessels **(B and D)** and keep the vessels separated **(C)**. The discriminating performance of our SBall filter is also significantly higher than Sato's, which detects many unwanted structures **(E)**. However, in at least one difficult location **(F)** Sato's filter resulted in improved connectivity at a junction. Overall, we found that the proposed filter is superior for small vessel detection, for keeping the structure clear, and for background discrimination.

4.4 Discussion and conclusion

The new curvilinear structure detection filter proposed in this paper was motivated by recent advances in HARDI image processing. By analyzing the values of a difference operator applied in many directions, our new filter is able to achieve a high level of discriminations while offering good performance at vessel junctions. The SBalls form the basis of the vessel detection filter are represented using spherical harmonics. This provides a natural way to interpolate the SBall value in any direction, a property that was, in fact, used in the definition of the vesselness measure $\mathcal{S}(\mathbf{x})$ in (4.11). At the same time, when combined with our contrast invariant




Figure 4.11 Vessel detection in a cropped clinical TOF-MRA volume. *Left column:* sagittal anterior view, and *right column:* lateral left view. *Top to bottom:* Proposed, Frangi's, and Sato's vesselness filter response. The arrows indicate some regions of interest. See the text for a discussion.

diffusivity index, the SH coefficient can be used to provide a rich visualization of the image structure, as in Fig. 4.2 and 4.6. We believe that their reasonable memory footprint will lead




Figure 4.12 Vessel detection in a cropped clinical cardiac CTA volume. *Left column:* axial superior view, and *right column:* sagittal anterior view. *Top to bottom:* vesselness response for the SBall filter, Frangi's filter, and Sato's filter. The arrows indicate some regions of interest. See the text for a discussion.

to applications where the SBall will provide directional information to a tracking or segmentation process. Moreover, our analysis of the SBall is not limited to the vesselness measure. Other quantities can be derived from it, such as the flux and the oriented flux. The proposed

SBall vesselness filter has also been integrated into the scale-space framework of Lindeberg (Lindeberg, 1996) for multiscale vessel detection.

Experiments on three different synthetic images demonstrate the performance of the SBall vessel detection filter. In the first, we showed that the multiscale integration of the filter provides an appropriate framework for detecting vessels within a reasonable range of size without favoring any particular scale. In the second and third experiments, we showed, both qualitatively and quantitatively, that the SBall filter provides the best balance between sensitivity and discriminative power. Comparisons are provided with the two classic filters of Frangi and Sato as well as with the newer Qian filter.

Finally, the performance of the SBall vessel detection filter has been assessed on two clinical medical images from different imaging modalities: TOF-MRA and CTA. In this case, the results were compared to those of Frangi's filter and Sato's filter. The findings in this section tends to corroborate those in the experiment with synthetic images. In summary, the SBall filter performs better than the other filters at preserving vessel junctions and small vessels, while having good resilience to the presence of unwanted structures. Building on these encouraging results, future work will be directed toward integration of the SBall into tracking and segmentation methods. We also encourage other researchers to test our code on their own datasets.

4.5 Acknowledgement

The authors acknowledge Dr. Nagib Dahdah, cardiologist at Ste-Justine Hospital for children in Montreal, Quebec, who provided the CTA scan for our research. In addition, the authors would like to thank NSERC of Canada and Quebec's FQRNT for their financial support.

CHAPTER 5

NON-RIGID 2D/3D REGISTRATION OF CORONARY ARTERY MODELS WITH LIVE FLUOROSCOPY FOR GUIDANCE OF CARDIAC INTERVENTIONS

David Rivest-Hénault¹, Hari Sundar² and Mohamed Cheriet¹,

¹ Département de génie de la production automatisée, École de Technologie Supérieure,
1100 Notre-Dame Ouest, Montréal, Québec, Canada H3C 1K3

² Imaging and Visualization, Siemens Corporate Research
755 College Road East, Princeton, NJ, 08540 U.S.A.

Paper accepted for publication in IEEE Transaction on Medical Imaging in April 2012

Abstract

A 2D/3D non-rigid registration method is proposed that brings a 3D centerline model of the coronary arteries into correspondence with bi-plane fluoroscopic angiograms. The registered model is overlaid on top of interventional angiograms to provide surgical assistance during image-guided chronic total occlusion procedures, thereby reducing the uncertainty inherent in 2D interventional images. The proposed methodology is divided into two parts: global structural alignment and local non-rigid registration. In both cases, vessel centerlines are automatically extracted from the 2D fluoroscopic images, and serve as the basis for the alignment and registration algorithms. In the first part, an energy minimization method is used to estimate a global affine transformation that aligns the centerline with the angiograms. The performance of nine general purpose optimizers has been assessed for this problem, and detailed results are presented. In the second part, a fully non-rigid registration method is proposed and used to compensate for any local shape discrepancy. This method is based on a variational framework, and uses a simultaneous matching and reconstruction process to compute a non-rigid registration. With a typical run time of less than 3 seconds, the algorithms are fast enough for interactive applications. Experiments on five different subjects are presented and show promising results.

Keywords: 2D/3D Registration, X-ray fluoroscopy, computed tomography, image-guided interventions, interventional cardiology, chronic total occlusions

5.1 Introduction

X-ray fluoroscopy is the modality of choice for the guidance of percutaneous coronary interventions (PCI) of chronic total occlusions (CTO). During these procedures, crossing CTOs using a guidewire is particularly hazardous, since the occlusion blocks the propagation of the contrast agent, and makes the occluded portion of the vessel invisible under fluoroscopy. In addition to the contrast issue, the projective nature of fluoroscopy results in ambiguities in the interpretation of 3D structures that further complicates the interventional procedure. Similar challenges arise in a vast array of minimally invasive procedures (Peters, 2006; Peters and Cleary, 2008).

Current pre-interventional planning routinely includes the acquisition of a computed tomography angiography (CTA), or other 3D imaging modality, which is used to reduce visual uncertainty. The performance characteristics of CTA with blood pool contrast injection enables the calcifications causing the CTOs to be clearly distinguished, as opposed to what can be perceived with interventional X-ray fluoroscopy with direct contrast agent injection. This makes CTA a highly valuable tool at the planning stage, but establishing correspondence between these data and the interventional images can prove difficult.

To address this issue, a 3D model can be extracted pre-procedure from the acquired volume, aligned with the 2D fluoroscopic views, and overlaid on top of the live images, thereby augmenting the interventional images. However, achieving this alignment is a challenge in itself, mainly because finding intermodal correspondence is a non trivial problem, and also because of the non linear aspect of the underlying optimization problem. Furthermore, a simple rigid transformation might not be sufficient to provide a satisfying 2D/3D registration. Since the 3D planning image is acquired under a breathhold, there are significant shape changes as compared to the intraoperative images acquired under free breathing. This makes it extremely important to use a non-rigid registration method while aligning the preoperative model with the live intraoperative images. This is, however, a difficult ill-posed inverse problem. Nevertheless, 2D/3D registration methods have the potential to greatly reduce the uncertainty relative to interventional X-ray angiography, and to do so with only minimal modification to the existing clinical flows. The development of such methods is the main objective of this paper, and its most significant contribution.

From a broader perspective, 2D/3D registration methods have numerous applications in fields such as neurology (Hipwell *et al.*, 2003), orthopaedics (Benameur *et al.*, 2003), and cardiology (Ruijters *et al.*, 2009). The associated body of literature is expanding rapidly, as is apparent in

the thorough review of techniques recently published by Markelj *et al.* (Markelj *et al.*, 2012). Below, the techniques most closely related to the one presented in this paper are discussed.

The maximal precision that can potentially be achieved by a registration process is directly linked to the complexity of the transformation model involved. Many 2D/3D registration approaches consider only a rigid transformations model (Feldmar *et al.*, 1997; Turgeon *et al.*, 2005; Truong *et al.*, 2009; Ruijters *et al.*, 2009), or a slightly more flexible affine model (Tsin *et al.*, 2009). In general, this is appropriate for rigid structures, such as bones, or to provide an initial alignment of the modalities, but might prove insufficient to account for the shape changes of flexible structures. Consequently non-rigid deformation models have been proposed for 2D/3D registration (Groher *et al.*, 2009; Metz *et al.*, 2009; Liao *et al.*, 2010). In (Groher *et al.*, 2009), the transformation model is strongly constrained by a length conservation term. The method has mostly been demonstrated on synthetic examples, and its computational complexity, resulting in a computational time of around 5 min, makes it impractical in an interventional setting. The approach by Metz *et al.* (Metz *et al.*, 2009) requires the acquisition of 4D CTA, that allows for the use of a 3D+t model for cardiac motion. However, it does not account for shape changes occurring through respiratory movement, and the acquisition of a 4D CTA is not possible in most interventions, which is a major shortcoming of this approach. Interesting results are presented in the paper by Liao *et al.* (Liao *et al.*, 2010), but their technique makes use of features specific to the abdominal aorta.

In a clinical setting, one or two fluoroscopic planes are used for intervention guidance. Some authors (Bouattour *et al.*, 2005; Groher *et al.*, 2009; Duong *et al.*, 2009) have investigated the monoplane scenario, but the reported errors are large in the out-of-plane direction. This lack of accuracy seriously limits confidence in the registration process. Biplane acquisition greatly reduces the ambiguity associated with these interventions, which is why this paper focuses on the alignment and registration of a 3D model with two fluoroscopic images. In our experience of five vastly different clinical sites located in Canada, Germany and the Netherlands, we have observed that biplane acquisitions are now performed on a daily basis for complex cardiac interventions. Thus, the proposed methodology is practical in many clinical situations.

The method presented in this paper is designed to non-rigidly align a preoperative 3D centerline model of the coronary arteries with two intraoperative fluoroscopic images to visually augment the interventional images. It is composed of two steps: 1) a global transformation model is calculated to provide an initial rigid or affine alignment, and 2) a fully non-rigid model is used to compute the final registration. To estimate the global alignment parameters, a formulation derived from (Sundar *et al.*, 2006) is used. This formulation benefits from distance maps

(Fabbri *et al.*, 2008; Paragios *et al.*, 2003) to measure the discrepancies between the projections of the 3D model and features automatically extracted from the 2D images. Such formulations have less computational complexity than those based on point matching approaches (Stewart, 2006; Truong *et al.*, 2009) or on digitally reconstructed radiographs (Prümmer *et al.*, 2006; Souha Aouadi, 2007). A contribution of this paper is to generalize the cost function defined in (Sundar *et al.*, 2006) to cover the affine and biplane case.

The optimization surface in 2D/3D registration problems can be highly non-linear for several reasons: the discretization of the image, the complexity of the structures to be registered, and the use of a rigid or affine transformation model. As a result, the minimization of the cost function is challenging. As we know that no single optimization method outperforms all the others, it is hard to understand why only one or two optimization algorithms have been evaluated in most related works (Turgeon *et al.*, 2005; Sundar *et al.*, 2006; Ruijters *et al.*, 2009). In addition, work published by Lau and Chung (Lau and Chung, 2006) suggests that global optimizers might perform better than the more popular local optimizers for a related 2D/3D registration problem. An important contribution of this work is thus to present a rigorous comparison between seven local and two global optimizers. The data were gathered using one realistic simulated case and five clinical cases; 2D and 3D errors are reported as well as runtime measurements. The results presented here can help implementers choose the best algorithm for their application. A similar study, but for intensity-based 2D/3D registration, is presented in (van der Bom *et al.*, 2011).

Work recently published by Ruijters *et al.* (Ruijters *et al.*, 2009) shares some similarities with the proposed 2D/3D alignment method. Their registration method makes use of the distance transform of the projection of a 3D centerline and of the output of the Frangi vesselness filter (Frangi *et al.*, 1998) to compute the cost associated with a certain pose. The optimization is carried out with either a Powell optimizer or a stochastic optimizer. There are, however, a few important differences between the approaches. While they specifically avoid 2D segmentation, we propose using a very recent 2D segmentation algorithm (Schneider and Sundar, 2010). By doing so, it is possible to precompute the distance transform on the 2D images, instead of computing the distance transform of the projection of the 3D centerline for a certain pose at each iteration. This can lead to an improvement of up to 2 orders of magnitude in registration time. Nonetheless we acknowledge that computing the automatic 2D segmentation incurs some small overhead. We also make use of the affine transformation model in addition to the rigid one, and present a more thorough evaluation using additional optimizers. Finally, (Ruijters *et al.*, 2009) does not consider non-rigid registration.

The 2D/3D registration method can also be used to capture the motion of the structure of interest across a sequence of frames (Blondel *et al.*, 2006; Tsin *et al.*, 2009). The demonstration of the suitability of the proposed global registration method in this multiframe setting is another contribution of this work. With respect to the non-rigid registration method, our main contribution is the formulation of this problem as a simultaneous matching, reconstruction and registration problem, that can, on modern hardware, be solved fast enough to be used intraoperatively during CTO procedures.

The rest of this paper is organized as follows. Section 5.2 presents the practical system considered here and other background information. The global alignment method is described in section 5.3. The proposed non-rigid registration method is introduced in section 5.4. Experiments with both the global alignment method and the non-rigid registration method are presented in section 5.5. Finally, a discussion and the conclusion are presented in section 5.7.

5.2 Background Information

A 3D centerline representation of a coronary artery tree, segmented (Gülsün and Tek, 2008) from a pre-operative CTA volume, is to be non-rigidly registered to 2 simultaneous fluoroscopic images. Starting from the default 3D location, computed from the calibration of the apparatus, the 3D alignment of the centerline is progressively refined using translation-only motions, a rigid transformation, and an affine transformation. Lastly, a non-rigid transformation is computed and provides the final registration.

The geometry of the system under consideration is represented schematically in Fig. 5.1 and is described using five coordinate systems (CS). The 3D centerline representing the coronary tree is described with respect to U , and a CS C is centered at the reference point of the 3D imaging device. The 3D centerline itself is described by using a set of S segments composed of Q_s control points $\mathbf{x}_{s,q} \in \mathbb{R}^3$ forming an undirected acyclic graph. In term of 2D angiography, X is the reference CS of the biplane C-arm, and CSs P_1 and P_2 are centered on the 2D imaging planes. The transformations between C and U , and between X and $\{P_1, P_2\}$, are known from the calibration of the apparatus, and are encoded in the rigid transformation matrices \mathbf{C} , \mathbf{P}_1 and \mathbf{P}_2 respectively. The projective geometry of the two X-ray planes is also assumed to be available, and is encoded in the projection operators Ψ_1 and Ψ_2 that map a 3D point in homogeneous coordinates $\mathbf{x} = [x, y, z, 1]^T$ to a 2D point $\mathbf{x}' = [x, y]$ on the corresponding fluoroscopic plane.

The vessels are automatically segmented from the input fluoroscopic images, using the method proposed in (Schneider and Sundar, 2010), and represented as a binarization $B_n(i, j)$ of the in-




Figure 5.1 Geometry of the imaging system. Labeled frames represent coordinate systems. Dashed lines and bold symbols represent transformation matrices.

put image. $B_n(i, j) = 1$, if the input image pixel at location $(i, j) \in \{0, \dots, W - 1\} \times \{0, \dots, H - 1\}$ corresponds to some structure of interest, and $B_n(i, j) = 0$ otherwise. It is assumed that the two fluoroscopic images are the same size $W \times H \in \mathbb{N}^2$. For the sake of completeness, the automatic segmentation method in (Schneider and Sundar, 2010) is briefly outlined here: 1) an input fluoroscopic image is processed using the Hessian-based Sato *vesselness* filter (Sato *et al.*, 1998); 2) seed points are generated by sparsely sampling a thresholding of the *vesselness* measure; 3) starting from those points, fibers are generated by integrating along the first eigenvector of the Hessian matrix computed at each pixel location; and 4) the fiber bundles are iteratively thinned to extract the centerlines of the vessels. The default parameters of this method have been used without modification throughout this paper. Examples of outputs are presented in Fig. 5.2¹.

¹All 3D centerline projections presented in this work are color-coded as follows: if the 3D centerline is locally parallel to the view plane, it is green; if it is perpendicular, it is red. A linear interpolation is used for intermediary situations.

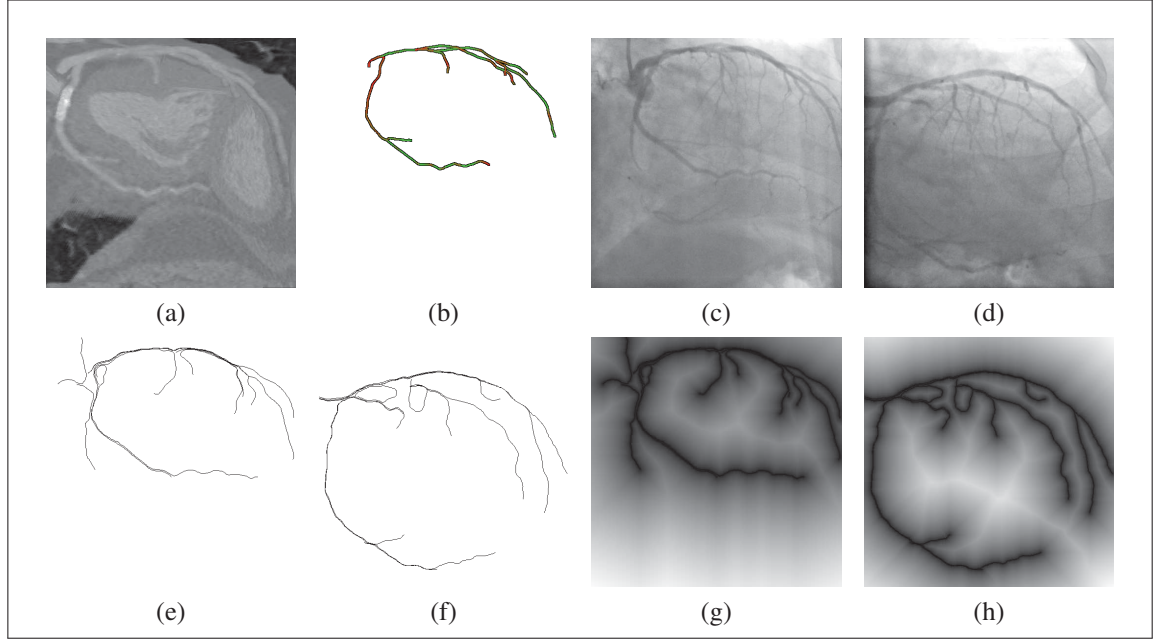


Figure 5.2 Input data. a) curved slice from a CT volume; b) 3D centerlines; c) and d) fluoroscopic angiograms from 2 different points of view; e) and f) segmentation of the fluoroscopic images; and g) and h) distance transforms of the segmentations.

5.3 Translational, rigid, and affine alignment

In a typical clinical setting, the transformation that aligns the 3D coronary centerlines from the CTA acquisition to the images of the same structures on the biplane fluoroscopy can only be measured approximately, since the position of the patient cannot be controlled with a high degree of accuracy. A transformation \mathbf{T} , that is a mapping between the two imaging modality coordinate systems, can be estimated by a calibration of the apparatus, and refined by taking into account the geometry of the structures of interest. Starting from the initial conditions, the improved alignment transformation is computed using a minimization process to estimate the parameters of a global translational or rigid transformation model. From that point, an affine transformation model can also be used to deform the 3D centerline in order to compensate for the shape discrepancy between the CTA and the biplane X-rays. A technique, inspired by (Sundar *et al.*, 2006), but extended to cover the biplane case and the use of affine transformations, is proposed for this purpose.

Formally, the total distance between the projection of the 3D centerline and the vessel centerlines segmented on each fluoroscopic plane is to be minimized. Let $\chi(\mathbf{x}) = \mathbf{T} \cdot \mathbf{C} \cdot \mathbf{x}$ be an affine transformation operator that maps a point $\mathbf{x} = [x, y, z, 1]^T$, relative to U , to the X coordinate system. Also, let $\Pi_n(\mathbf{x}) = \Psi_n \cdot \mathbf{P}_n \cdot \mathbf{x}$ be a projection that maps a point from X to the fluoroscopic

plane P_n . We define the energy

$$E_{\text{Global}}(\chi) = \sum_{n=1}^2 \sum_{s,q} D_{B_n}(\Pi_n(\chi(\mathbf{x}_{s,q}))) \quad (5.1)$$

as the quantity that needs to be minimized. Here, $D_{B_n}(x)$ is the distance between a 2D point and the closest point where $B_n(i, j) = 1$. The distance $D_{B_n}(x)$ for each image position is computed beforehand, by creating the distance transform of each $B_n(i, j)$ segmentation, as presented in Fig. 5.2.

Optimizing the 12 values of \mathbf{T} directly is both ineffective and inefficient, and so a parametric approach has been used. The translation-only, rigid, or affine transformations \mathbf{T}_T , \mathbf{T}_R , and \mathbf{T}_A are represented using the parameter sets $\mathcal{T} \in \mathbb{R}^3$, $\mathcal{R} \in \mathbb{R}^6$, $\mathcal{A} = \mathbb{R}^{12}$ respectively. The mappings from the parametric to the matrix representations of the \mathbf{T} s are defined as follows (Shoemake and Duff, 1992):

$$\begin{aligned} \mathbf{T}_T(\mathcal{T}) &= M_T(\mathcal{T}_1^3) \quad , \\ \mathbf{T}_R(\mathcal{R}) &= M_T(\mathcal{R}_1^3) \cdot M_R(\mathcal{R}_4^6) \quad , \text{ and} \\ \mathbf{T}_A(\mathcal{A}) &= M_T(\mathcal{A}_1^3) \cdot M_R(\mathcal{A}_4^6) \\ &\quad \cdot M_R(\mathcal{A}_{10}^{12}) \cdot M_S(\mathcal{A}_7^9) \cdot (M_R(\mathcal{A}_{10}^{12}))^T \quad , \end{aligned} \quad (5.2)$$

where \mathcal{S}_i^j is the subset containing the elements with indices $\{i, \dots, j\}$ of \mathcal{S} and

$$\begin{aligned} M_T(\mathcal{T}) &= \begin{bmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}, M_S(\mathcal{S}) = \begin{bmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \\ \text{and} & \\ M_R(\theta) &= \begin{bmatrix} c_y c_z, & c_z s_x s_y - c_x s_z, & s_x s_z + c_x c_z s_y, & 0 \\ c_y s_z, & s_x s_y s_z + c_x c_z, & c_x s_y s_z - c_z s_x, & 0 \\ -s_y, & c_y s_x, & c_x c_y, & 0 \\ 0, & 0, & 0, & 1 \end{bmatrix}, \end{aligned} \quad (5.3)$$

where $c_x = \cos(\theta_x)$ and $s_x = \sin(\theta_x)$. The transformation matrices \mathbf{T}_T , \mathbf{T}_R , and \mathbf{T}_A are re-centered in such a way that any change in the rotation parameters induces a motion that appears to occur around the origin of the U coordinate system, as opposed to the origin of C .

This operation does not influence the energy function (5.1), but is usually beneficial for the optimizers. Intuitively, this is because such re-centering reduces the distance, in parameter space, associated with a rotation or a scaling operation that appears to occur at the centroid of the CT volume. The final definition of \mathbf{T} is then

$$\mathbf{T} = \mathbf{C} \cdot \mathbf{T}_J \cdot \mathbf{C}^{-1} , \quad (5.4)$$

where \mathbf{C} is the transformation from U to C , and $J \in \{T, R, A\}$, depending on the transformation model chosen.

Any non-derivative general purpose optimization algorithm can be used to minimize (5.1). In this work, 7 local (L) and 2 global (G) algorithms have been tested for this purpose. These algorithms are briefly described here:

Best Neighbor (L): At each iteration, a step in all principal directions of the parameter space is considered. The best move is applied, and a new iteration begins. When no move improves the solution, the step is halved.

Nelder-Mead² (L): A classic numerical optimization method that minimizes an m -dimensional function by evaluating the function value at the $m + 1$ vertices of a general simplex (or polytope). At each iteration, the vertex with the worst value is replaced by another one using a reflection operation followed by either an expansion or a contraction operation (Nelder and Mead, 1965).

Sbplx³ (L): This method decomposes the problem in a low-dimensional subspace, and uses the Nelder-Mead algorithm to perform the search (Rowan, 1990).

Cobyla (L): The *Constrained Optimization BY Linear Approximation* method works by constructing linear-approximations of the cost function and constraints using $m + 1$ vertex simplexes, and minimizes this approximation. The radius of the simplex is progressively reduced, while maintaining a regular shape (Powell, 1994, 1998).

Bobyqa (L): Each iteration of the *Bound Optimization BY Quadratic Approximation* method use a quadratic approximation of the cost function F constructed, typically by considering

²Also known as *Downhill Simplex*.

³The name of the original author's implementation is *Subplex*, the name of the NLOpt (Johnson, 2012) implementation is *Sbplx*.

$2m + 1$ interpolation point. The trust region is progressively reduced until there is no further improvement (Powell, 2009).

Powell-Brent (L): At each iteration, a succession of exact 1D line optimizations is performed using Brent's method. The solution is updated using Powell's method of conjugate search directions (Powell, 1964; Brent, 1973).

Praxis (L): Brent's *PR*incipal-*AXIS* method is a refinement of Powell's method of conjugate search directions (Brent, 1973).

Differential Evolution (G): This is a population-based stochastic global optimization method, the main feature of which is the use of the vector of the difference between pairs of individuals as the basis for the population evolution (Storn and Price, 1997).

Direct (G): This global optimization algorithm is designed for problems with finite bound constraints, as is the case here. The parameter space is systematically and deterministically searched by dividing it into smaller and smaller hyperrectangles (Jones *et al.*, 1993).

The free parameters of each algorithm have been adjusted in accordance with the recommendations of their original author or implementer. The implementations of the following algorithms are taken from the NLOpt library (Johnson, 2012): *Praxis*, *SBPLX*, *Cobyla*, *Bobyqa*, and *Direct*. The performance of these optimizers is discussed in detail in section 5.5. In order not to change the nature of the problem, generous bounds have been fed to the optimizers: $\pm 200\text{mm}$ for the translation parameters, $\pm 45\text{deg}$ for the rotation parameters, $[0.5, 1.5]$ for the scale parameters, and $\pm 360\text{deg}$ for the scale-rotation parameters (\mathcal{A}_{10}^{12} in (5.2)).

The same set of stopping criteria was used for all the optimizers, except *Differential Evolution*. Specifically, the optimizers were allowed to run until, after one optimization step: 1) the value of the cost function was reduced by less than $1e^{-14}$; or 2) the value of the optimized parameter changed less than $1e^{-4}\text{mm}$ for the translation parameters, $1e^{-5}\text{rad}$ for the rotation parameters, and $1e^{-7}$ for the scaling parameters. The parameter tolerances have an intuitive interpretation. For example, a difference in translation of $1e^{-4}\text{mm}$ corresponds to a maximal displacement of $1/2000$ of a pixel on an image plane. As to the *Differential Evolution* optimizer, since the convergence of stochastic algorithms is not regular, we decided to let the optimizer run until it had evaluated the cost function 100,000 times⁴, that corresponds to a median run of approximately 10s.

⁴1,000,000 times, in the case of simulation.

Starting from the default patient position, the alignment is progressively refined by considering 1) translations only, 2) rigid transformations, and 3) affine transformations.

5.3.1 Multi frame alignment

The global alignment method can be expanded to cover the multi-frame scenario. In this case, the crucial hypothesis is that because of the temporal continuity, the estimated alignment transformation would only change little from one frame to the other. The global energy function can then be re-defined for a sequence of F biplane frames:

$$E_{\text{Multi}}(\chi_f) = \sum_{f=1}^F \sum_{n=1}^2 \sum_{s,q} D_n(\Pi_n(\chi_f(\mathbf{x}_{s,q}))) + \gamma \sum_{f=1}^{F-1} D(\chi_f, \chi_{f+1}) , \quad (5.5)$$

where γ is a free parameter, and χ_f is defined in a way similar to χ , but with a different transformation matrix \mathbf{T}_f for each time point. The most delicate part of this energy function is the definition of the inter-frame distance $D(\chi_f, \chi_{f+1})$ since it is well known that the space of rigid, and affine, transformation matrices is not linear. The rigid transformation space can be locally linearized for small changes in orientation (Boisvert *et al.*, 2008; Pennec and Thirion, 1997), as is assumed to be the case here. Using this framework, a rigid transformation is represented by a translation vector \vec{t} and a rotation vector \vec{r} , defined as the product of a unit vector \vec{n} and a rotation angle θ , such that $\vec{r} = \theta \vec{n}$. The translation vector simply corresponds to the translational part of the transformation matrix, and the mapping between the rotation matrix R and the rotation vector \vec{r} is given by (Boisvert *et al.*, 2008):

$$\theta = \arccos\left(\frac{\text{Trace}(R) - 1}{2}\right) \quad \text{and} \quad \begin{pmatrix} 0 & -n_z & n_y \\ n_z & 0 & -n_x \\ -n_y & n_x & 0 \end{pmatrix} = \frac{R - R^T}{2 \sin(\theta)} .$$

Using this representation, the inter-frame distance is defined as follows (Boisvert *et al.*, 2008):

$$D(\chi_f, \chi_{f+1}) = N_\alpha(\mathbf{T}_{f+1}^{-1} \cdot \mathbf{T}_f) . \quad (5.6)$$

with

$$N_{\alpha}(R) = \|\vec{r}\|^2 + \|\alpha\vec{t}\|^2 . \quad (5.7)$$

Here, the transformations \mathbf{T}_f have the same meaning as the transformations \mathbf{T} defined in the previous section. Also, α is a real number that balances the respective contributions of the rotation and translation parts. It was set to 0.05, in accordance with the recommendation in (Boisvert *et al.*, 2008).

Defining a linear distance between affine transformation matrices is more involved and requires complex computation (Arsigny *et al.*, 2006). Since it is assumed in this work that rigid transformations capture the major part of the transformation, we decided to regularize only the rigid part of the transformation and leave the remaining affine parameters unconstrained. Formally, in the case of affine transformations, the transformations \mathbf{T}_f in the regularizer definition are composed using the parameters \mathcal{A}_1^6 only, whereas the full set of 12 parameters is used to compose the transformations χ_f in the energy definition.

This energy function can be useful for tracking the alignment transformation over a sequence of frames, or, if $\gamma \rightarrow \infty$, for computing an average transformation using, for example, three adjacent frames. As with the single frame energy, the transformations χ_f are progressively refined by sequentially increasing the complexity of the transformation models.

5.4 Non-rigid registration

Global affine transformation models cannot entirely compensate for the inter-modal shape discrepancy caused by breathing and by the beating of the heart. A non-rigid registration method is now introduced, that can greatly improve the visual correspondence between the 3D coronary tree centerlines and the two calibrated fluoroscopic images. The aim of this method is to be automatic, in the sense that it does not require the user to identify correspondences.

The non-rigid transformation is represented as a set \mathbf{r} of 3D translation vectors $\mathbf{r}_{s,q}$ that are to be applied to the corresponding centerline points $\mathbf{x}_{s,q}$ in CS U . Thus, a registered point $\bar{\mathbf{x}}_{s,p}$ is computed using $\bar{\mathbf{x}}_{s,q} = \mathbf{x}_{s,q} + \mathbf{r}_{s,q}$. An energy function is defined to measure the quality of a solution:

$$E_{\text{NR}}(\mathbf{r}) = E_{\text{Image}}(\mathbf{r}) + E_{\text{Internal}}(\mathbf{r}) . \quad (5.8)$$

An iterative minimization process is used to minimize (5.8), and, starting from the position $\mathbf{r}_{s,q} = [0, 0, 0]^T \mid \forall \{s, q\}$, the transformations $\mathbf{r}_{s,q}$ are progressively refined by following a gradient descent approach. This results in a smooth motion of the registered centerline, as can be seen in Fig 5.3.




Figure 5.3 Progression of a non-rigid registration. From left to right: sample input image, and position of the centerline after 0, 50, 350, and 1200 iterations. This corresponds to subject 4 in Table 5.3. *See also the related video file, as described in section 5.6.*

5.4.1 Image energy

A matching and reconstruction process is used to calculate a reconstructed point $\hat{\mathbf{x}}_{s,q}$ for every point $\mathbf{x}_{s,q}$ in the 3D centerline model. The matching is performed on each fluoroscopic plane separately. Starting from the projection \mathbf{x}' of $\bar{\mathbf{x}}_{s,q}$ on an image plane n , a search on the corresponding segmentation image B_n is performed in the directions that are perpendicular to the projection of the centerline – see Fig. 5.4-a) for an illustration. This direction is computed, at every location, based on the partial derivatives $\left(\frac{\partial \mathbf{x}'}{\partial x}, \frac{\partial \mathbf{x}'}{\partial y}\right)$. As a result, at most two matching points, $\mathbf{x}''_{n,a}$ and $\mathbf{x}''_{n,b}$, are identified on each fluoroscopic plane $n \in \{1, 2\}$. The pair $(\mathbf{x}''_{1,i}, \mathbf{x}''_{2,i})$, $i \in a, b$ that best satisfies the epipolar constraint is kept. For this purpose, the following quantity is considered:

$$d(\mathbf{x}''_{1,i}, \mathbf{x}''_{2,i}) = d(\mathbf{x}''_{1,i}, e_{2,i}) + d(\mathbf{x}''_{2,i}, e_{1,i}) \quad , \quad (5.9)$$

where $d(\mathbf{x}''_{n,i}, e_{m,i})$ is the Euclidean distance between point $\mathbf{x}''_{n,i}$ and the epipolar line induced by $\mathbf{x}''_{m,i}$, $\{n, m\} \in \{1, 2\}$, with $n \neq m$. The pair minimizing (5.9) is then used to calculate a reconstructed 3D point $\hat{\mathbf{x}}_{s,p}$ using the projective geometry of the system (Hartley and Zisserman, 2004). If only one point $\mathbf{x}''_{n,i}$ is found on an image plane, it is retained by default.

Searching for matching points in directions that are locally perpendicular to the projection of the centerline is an important aspect of this process. Since it is known *a priori* that the

global alignment of the structures is correct, the perpendicular search direction prevents many unrealistic matches, compared to a nearest-point approach. This is especially important when the 2D automatic segmentation is not perfect. For example, when the 3D projection results in a curve that is longer than the 2D curve that has been automatically segmented, taking the nearest point would result in a significant shortening of the curve from its tip. In contrast, considering the perpendicular direction considerably reduces this effect, because no match would be found at the tip of the 3D centerline projection.

The quality of the reconstruction cannot be guaranteed and an indicator function $H(\mathbf{x}_{s,p}) \mapsto \{0, 1\}$ is used to reject outliers. $H(\mathbf{x}_{s,p}) = 0$ if: 1) no matching point is found in any image planes; 2) the 2D distance between \mathbf{x}' and \mathbf{x}''_n is greater than Max_{2D} ; and 3) the 3D distance between $\mathbf{x}_{s,p}$ and $\hat{\mathbf{x}}_{s,p}$ is greater than Max_{3D} . Otherwise, $H(\mathbf{x}_{s,p}) = 1$. Here, Max_{2D} and Max_{3D} are thresholding functions. More complex kernels, such as Huber functions, could be integrated in the proposed scheme, but this possibility has not been explored. Their values were fixed to $\text{Max}_{2D} = 50$ pixels and $\text{Max}_{3D} = 6\text{mm}$.

The image energy in (5.8) is defined as the sum of the squared error between the reconstructed points $\hat{\mathbf{x}}_{s,p}$ and the centerline point, for all valid points, as follows:

$$\begin{aligned} E_{\text{Image}}(\mathbf{r}) &= \sum_{s,q} H(\mathbf{x}_{s,q}) (\hat{\mathbf{x}}_{s,q} - (\mathbf{x}_{s,q} + \mathbf{r}_{s,q}))^2 \\ &= \sum_{s,q} H(\mathbf{x}_{s,q}) (\hat{\mathbf{x}}_{s,q} - \bar{\mathbf{x}}_{s,q})^2. \end{aligned} \quad (5.10)$$

This energy is at its minimum when all the registered centerline points are at the same position as the reconstructed points.




Figure 5.4 a) Matching of a point from the centerline x' with a point from the 2D structure. Points x''_a and x''_b are identified as potential candidates. b) Creation of 3 myocardium constraints (k_1 , k_2 , and k_3) at a junction.

5.4.2 Internal energy

Regularizers are used to keep the registered centerline visually coherent and geometrically plausible. The following three internal energy terms are considered for this purpose:

$$\begin{aligned}
 E_{\text{Internal}} &= E_{\text{Disp}}(\mathbf{r}) + E_{\text{Smooth}}(\mathbf{r}) + E_{\text{Myocard}}(\mathbf{r}) \\
 &= \mu \sum_{s,q} |\mathbf{r}_{s,q}|^2 + \nu \sum_{s,q} |\dot{\mathbf{r}}_{s,q}|^2 + \lambda \sum_{\{i,j\} \in \mathcal{K}} \left(\frac{|\mathbf{r}_i - \mathbf{r}_j|}{|\mathbf{x}_i - \mathbf{x}_j|} \right)^2,
 \end{aligned} \tag{5.11}$$

where μ , ν , and λ are energy-balancing free parameters. The first term, $E_{\text{Disp}}(\mathbf{r})$, is minimal when the displacement owing to the non-rigid transformations is small. This regularizer is necessary to minimize 3D out-of-plane motions that are only poorly constrained by a pair of 2D angiograms. The second, $E_{\text{Smooth}}(\mathbf{r})$, is used to ensure smoothness over each vessel segment. In this section, $\dot{\mathbf{r}}_{s,q}$ and $\ddot{\mathbf{r}}_{s,q}$ represent, respectively, the first and second derivatives of the translation vector \mathbf{r} for the position (s, q) with respect to the segment parameter. At segment junctions, Neumann boundary conditions, $\dot{\mathbf{r}}_{s,q} = [0, 0, 0]^T$, are assumed.

With only the first two internal energy terms, the centerline model is flexible, and the motions of the vessels are independent of each other, except at junction nodes. In reality, the vessel motions are mechanically constrained not only by their own specific rigidity, but also because they are attached to the myocardium. In this respect, $E_{\text{Myocard}}(\mathbf{r})$ is intended to act as a minimalist model of the myocardium constraint, and is used to ensure a certain degree of rigidity at the vessel branches as it prevents small segments from collapsing onto bigger ones. This constraint is modeled by using artificial links around the junction of three segments, as depicted in Fig. 5.4-b). Let D_s be a distance parameter, with its value set to three times the average diameter of the vessels present at the junction. The node on each segments that is at a distance D_s from the junction is joined to the corresponding node of the other segments. In a way that is analogous to a mechanical spring, with each link generating a force that is proportional to its displacement ratio if it is either expanded or compressed during the registration process. In (5.11), \mathcal{K} is defined as the set of all pairs of node $\{i, j\}$ that have been linked together.

5.4.3 Energy minimization

The complete energy function is defined using (5.8), (5.10), and (5.11), and it is minimized by computing its Euler-Lagrange equations and following a gradient descent approach. Starting

from the point where $\mathbf{r}_{s,p} = [0, 0, 0]^T \mid \forall \{s, p\}$ at iteration $z = 0$, the resulting discretized update equation is

$$\mathbf{r}_{s,q}^{z+1} = \mathbf{r}_{s,q}^z + \varepsilon \left[H(x_{s,q})(\hat{\mathbf{x}}_{s,q}^z - \bar{\mathbf{x}}_{s,q}^z) - \mu \mathbf{r}_{s,q}^z + \nu \ddot{\mathbf{r}}_{s,q}^z + \lambda \sum_{\{i\} \in \mathcal{K}_{s,q}} (\mathbf{r}_i^z - \mathbf{r}_{s,q}^z) \right], \quad (5.12)$$

where the numerical step ε is a small constant and $\mathcal{K}_{s,q} \subset \mathcal{K}$ is the possibly empty set of constraints acting on node (s, q) . Equation (5.12) is evaluated iteratively until the total displacement stagnates. More formally, the total displacement is denoted $\Delta^z \equiv \sum_{s,q} |\mathbf{r}_{s,q}^z|$, and the stopping criterion is defined, for $z > \beta$, as $\Delta^z - \Delta^{z-\beta} < \zeta \Delta^\beta$. An example of the progression of Δ^z is shown in Fig. 5.5. As we can deduce from the shape of the curve, the final result is not very sensitive to the specific values of β and ζ . In practice, we used $\beta = 100$ and $\zeta = 0.001$.




Figure 5.5 Total displacement of the centerline as a function of the number of iterations — see section 5.4.3 for details.

5.4.4 Parameter selection

The three parameters μ , ν , and λ govern the behavior of (5.12) and must be selected for a specific task. Starting from a certain parameter set ($\mu = 0.10$, $\nu = 10.0$, $\lambda = 1.0$), some insight about the behavior of method can be acquired by varying each value individually, and by computing the mean 2D projection error and the 3D error (described in section 5.5). This methodology has been applied to the patient 4 of the clinical dataset, following an affine align-

ment, and also to a dataset with synthetic deformations (see section 5.5.1). The former depicts a RCA, and the latter, a LCA. The results are presented in Fig. 5.7.




Figure 5.6 Illustration of the behavior of parameter λ in the non-rigid registration method: (left) $\lambda = 0$, (right) $\lambda = 0.3$.

Parameter μ constrains the overall displacement, and the 2D error curves suggest that the computed error increases with its value. However, the 3D error curve clearly indicates a minimum when μ has a value in the 0.05–0.10 range. In addition, we found that the convergence rate increases with μ . For example, the registration used to produce the graph in Fig. 5.7-(top left) converged in 4880, 3320, and 1100 iterations with $\mu = \{0.001, 0.01, 0.1\}$, respectively. Parameter ν controls the rigidity of the model. In this case, the 2D projection error suggests using a value in the 0.50–2.00 range. However, the 3D error curve shows that higher values might result in more accurate registration. We thus choose to use a values of $\nu = 10.0$, which appears to be a good trade-off between the 2D and 3D error curves. Parameter λ balances the myocardium constraint, and helps to preserve the general integrity of the global shape. With the clinical RCA dataset, it was found that high value can overly constrain the motion, which results in increased errors. However, with the simulated LCA dataset, the situation appears almost inversed, as λ values of up to 10.0 are clearly beneficial. We believe that this situation can be explained by the differences in the nature of the deformations that affects the LCA and RCA datasets. A value $\lambda = 0.5$ was selected because it represents a good trade-off in the error measures, and because it improves the visual appearance in many cases, as depicted in Fig. 5.6. It might be better to use a different value for LCA and RCA datasets, but this idea was not pur-

sued here. Finally, the parameters were thus fixed to the following values in all subsequent experiments: $\mu = 0.05$, $\nu = 10.0$, and $\lambda = 0.5$, unless stated otherwise.

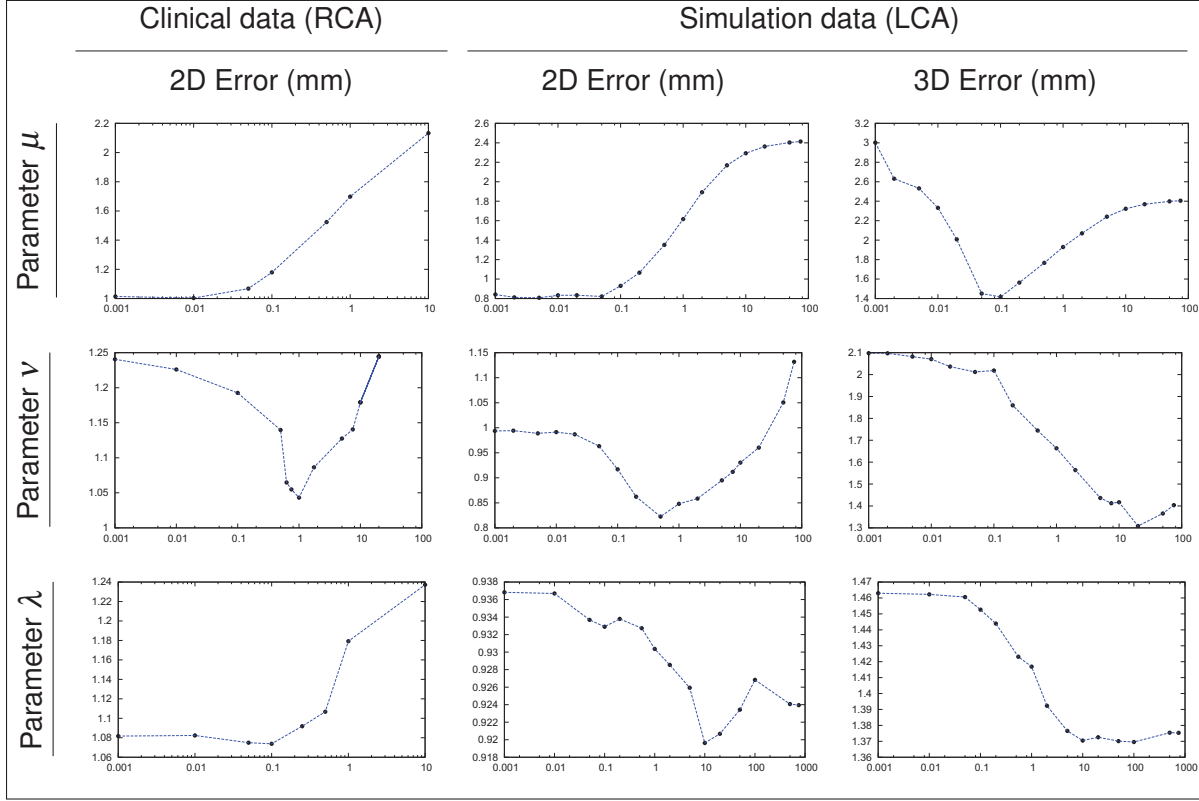


Figure 5.7 Mean 2D and 3D error after non-rigid registration. For each graph, one of the parameters was varied, while the others two were kept fixed. The default parameter values are: $\mu = 0.1$, $\nu = 10.0$, and $\lambda = 1.0$. The x-axis scales are logarithmic.

5.5 Experimental results

The algorithms were implemented in C++ within a proprietary prototyping environment that allows us to show an overlay of the 3D centerline model on the live fluoroscopy images during interventions. All global optimizers are single threaded. The non-rigid registration method is partly multithreaded using OpenMP. Run time given are for an i7 Q820 Quad Core Intel CPU. Presented run times don't account for 2D image segmentation. On average, this operation has a run time 0.200s per image, with a standard deviation of 0.088s. Experiments on both simulated and clinical data are presented.

5.5.1 Simulations

The proposed alignment and registration methods have been tested on a set of simulations. This has allowed us to characterize the performance of the various algorithms in a somewhat ideal scenario, and also to evaluate the 3D registration error, which is not possible using clinical data. The centerline was registered with a pair of digitally reconstructed radiographs (DRR), of size 512×512 pixels with a resolution of 0.345 pixel/mm, that were generated as follows: 1) the coronary arteries segmented in the CT volume were projected onto the simulated fluoroscopic planes using the technique presented in (Julien Couet and *et al.*, 2012), and 2) the inverse of the resulting projections was subtracted from the reference background images. This process is illustrated in Fig. 5.8, and is similar to that proposed in (Turgeon *et al.*, 2005). The simulations were used in the three scenarios described below. Since all the transformations are known during the simulation, all the 3D errors presented in this section are exact point-to-point errors. It should be noted that this style of error will penalize the compression or expansion of a segment, even when this produces no visual effect. It is thus more strict than most TRE formulations. The 2D error is computed as described in section 5.5.2, using the projection of the true centerline as the reference segmentation.

5.5.1.1 Dependence on the initial solution

The centerline is registered with a pair of DRRs, starting from different initial positions. This serves to quantify the sensitivity of the global alignment methods to a perturbation of the initial position, and equally, to a miscalibration of the \mathbf{C} , \mathbf{P}_1 , \mathbf{P}_2 or \mathbf{T} rigid transformation matrices. The selected initial positions are as follows: Let \mathbf{T}_0 be the transformation given by the calibration of the apparatus. Then, the 12 initial points are \mathbf{T}_0 displaced by $\pm\alpha$ mm along the three principal axes, and \mathbf{T}_0 rotated by $\pm\beta$ deg around the three principal axes. We used: $\alpha \in 10 * \{0, 1, \dots, 6\}$ mm, and $\beta \in 3.75 * \{0, 1, \dots, 6\}$ deg. The RMS error was computed for all results with the same level of perturbation (translation or rotation). The results depicted in Fig. 5.9 show that all algorithms presented are satisfactory when the rotational perturbation is less than 10 deg. Except for the *Neder-Mead* and the *Differential Evolution* algorithms, the error level increases rapidly past this threshold. All the algorithms in Fig. 5.9 seem to be able to recover from perturbations of up to 30mm in translation. At higher level, only the *Differential Evolution* and the *Best Neighbor* algorithms give consistent results, although the latter is less accurate. The *Neder-Mead* algorithm is more accurate and has a capture range of 50mm, which makes it a practical choice. The error curves for the *Cobyla*, *Bobyqa*, *Praxis*, and *Direct* algo-




Figure 5.8 Creation of the DRRs: The coronaries segmented in the CT volume are projected onto a simulated radiographic plane (a); this image is then subtracted from a real contrast-free angiography (b) to produce the final DRR (c). In one experiment, Gaussian noise is added (d), here with $\sigma = 50$.

rhythms were removed, because they were significantly higher, which impaired the interpretation of the graphs.

5.5.1.2 Robustness to image noise

Various amounts of Gaussian noise with variance $\sigma \in \{5, 10, \dots, 60\}$ were added to the fluoroscopic images, in order to assess the impact of image quality on the performance of the method. Each optimizer was run from 13 initial positions for each noise level (as above, with $\alpha = 5\text{mm}$,




Figure 5.9 Performance of the optimizers with respect to a perturbation of the initial position. *Bottom row*: missing points represent RMS error > 3 mm.

and $\beta = 7.5\text{deg}$, plus T_0). For each noise level, the RMS errors are presented in Fig. 5.10. The results of this test tend to demonstrate that the proposed methodology is not really sensitive to the presence of Gaussian noise in the fluoroscopies. The 2D and 3D error level shows very little correlation with the amount of added noise, even at levels as high as $\sigma = 60$ that are seldom observed in a clinical setting. This is an indication of the good performance of the selected 2D segmentation algorithm. In this test, the *Neder-Mead* and the *Differential Evolution* algorithms gave the most accurate and most consistent results. Again, the error curves for the *Cobyla*, *Bobyqa*, *Praxis*, and *Direct* algorithms were removed, because they were significantly higher.

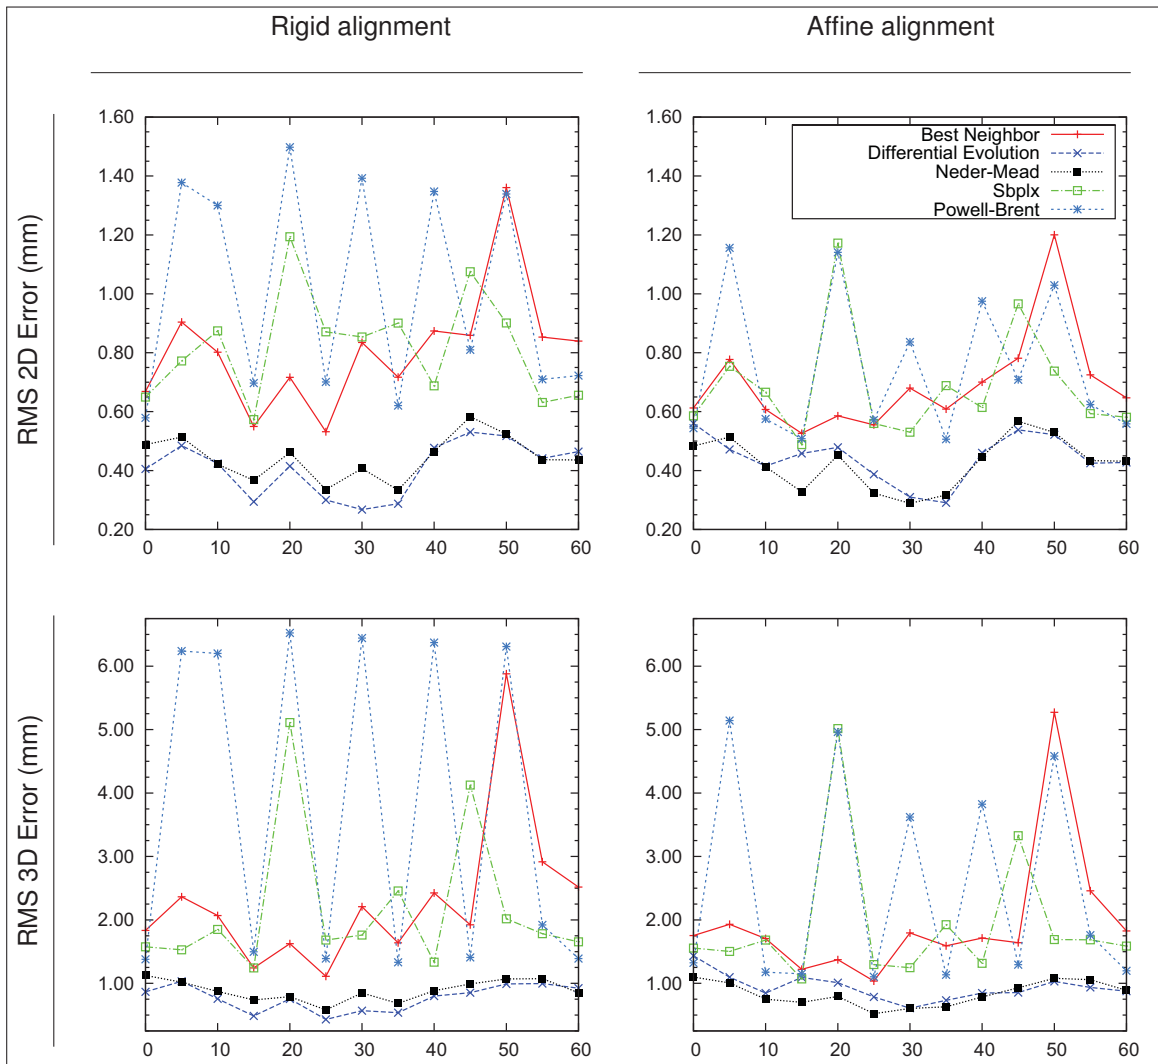


Figure 5.10 Performance of the optimizers with respect to the standard deviation of the input DRR noise. Gaussian white noise model with $\sigma \in [0, 60]$.

5.5.1.3 Non-rigid deformation

In these simulations, the centerline was deformed using a $3 \times 3 \times 3$ node thin plate spline (TPS) deformation model (Bookstein, 1989) covering the CT-scanned region. Each TPS node is moved toward the center node by a factor depending on the deformation parameter ξ . The 9 left nodes were shifted by a factor of ξ , the 9 center nodes, by 0.5ξ , and the 9 right nodes by 2ξ . Sample images are shown in Fig. 5.11. The non-rigid registration algorithm was tested for deformation levels $\xi \in \{0.020, 0.025, \dots, 0.060\}$, which resulted in mean and maximum 3D displacements of the centerline of $[1.209 - 3.638]$ mm and $[2.472 - 7.416]$ mm respectively. The nondeformed centerline curve was used as the initial solution, and no global alignment

was performed before applying the non-rigid registration algorithm. A quantitative evaluation is presented in Fig. 5.12. It was found that this algorithm gives good results up to the level $\xi = 0.045$. At this point, it reduces the average 3D error from 2.721mm to 1.198mm, which is qualitatively significant (Fig. 5.11). As will be demonstrated in the next sections, this level of performance is appropriate in a clinical scenario.




Figure 5.11 Sample non-rigid registration with simulated data. Top row: initial position; bottom row: final position.

5.5.2 Clinical data

Five datasets were used for the experiments presented in the following sections. Each dataset includes one CTA scan acquired at end-diastole (datasets 1 and 3) or end-systole (datasets 2, 4, and 5), and one biplane X-ray fluoroscopy recording. The two modalities are temporally




Figure 5.12 Residual 3D error with respect to the simulated non-rigid deformation level. *Baseline* correspond to the initial solution error.

aligned using ECG gating, and the coronary arteries were semi-automatically segmented in the CTA by a specialist using the technique described in (Gülsün and Tek, 2008). The mean 3D inter-point distances of the segmented centerlines were in the $[1.36 - 1.60]$ mm range. The angiograms image the left coronary arteries (LCA) in three cases and the right coronary artery (RCA) in the other two. The image size is 512×512 in all cases, and the image resolution is $\{0.345, 0.279, 0.279, 0.216, 0.279\}$ mm/pixel for dataset 1 to 5, respectively. The acquisition rate is 15 fps. Standard Siemens C-Arm calibration matrices were used. Although this is not a requirement of the method, matrices \mathbf{P}_1 and \mathbf{P}_2 were kept constant during acquisition (they are different for each dataset). The CTA and fluoroscopy exams had been prescribed to the patient for the treatment of a coronary disease.

The assessment of the performance of the alignment and registration algorithm was quantified using the mean 2D projection error.

$$Err(\chi, \bar{\mathbf{x}}_{s,q}) = \frac{1}{2M} \sum_{n=1}^2 \sum_{s,q} D_{S_n}(\Pi_n(\chi(\mathbf{x}_{s,q}))) , \quad (5.13)$$

where M is the total number of points in the 3D centerline model. Also, D_{S_n} is the distance transform of the reference 2D segmentation S_n , which was obtained by manual tracing on top of the corresponding fluoroscopy image.

5.5.3 Global alignment: evaluation of the performance of the optimizers

The performance of the nine optimizers in minimizing the global alignment energy (5.1) and in producing good quality 2D/3D alignment, were assessed using the following experimental

setup. For each of the five patient datasets, three pairs of fluoroscopic images were considered: the images temporally aligned with the CTA acquisition, and the previous and next adjacent frames. Temporal alignment helps to minimize the observable differences between the two modalities, thereby reducing the risk of failing a registration. This is reasonable clinically, as it would allow the registration matrix to be updated approximately once every second, in order to improve surgical guidance. Also, to test the robustness of the optimizers, 13 different initial points were used for the initialization. Let \mathbf{T}_0 be the transformation given by the calibration of the apparatus. The 13 initial points are \mathbf{T}_0 , \mathbf{T}_0 displaced by $\pm 5\text{mm}$ along the three principal axes, and \mathbf{T}_0 rotated by $\pm 7.5\text{deg}$ around the three principal axes. A total of $5 * 3 * 13 = 195$ experiments were thus conducted for each optimizer tested. As discussed in section 5.3, (5.1) is minimized by successively using: 1) translation-only transformation, 2) rigid transformation, and 3) affine transformation.

The mean residual energies left after minimizing with each of the three transformation models are presented in Table 5.1, along with the measured mean 2D projection error. In addition, per patient box-and-whisker plots of the mean 2D error are displayed in Fig. 5.14, and sample results are shown in Fig. 5.13.

The total computational time needed to successively estimate the translational, rigid, and affine alignments is reasonable for all local optimizers with median values of 105 ms or less, as can be seen in Fig 5.15. In all cases, the total time was under 1 s. The method is thus suitable for an interactive application with any of the local optimizers since the alignment appears to be computed almost instantaneously at the push of a button. In addition, the very short computational time of the *Neder-Mead* optimizer allows us to envision real time application. The computational time is much higher with the two global optimizers, as presented in Table 5.2. This means that they are only usable in an offline low-interaction setting.

A look at the average values of Table 5.1 reveals that the best algorithms for minimizing the energy function with the affine transformation model are the following: *Differential Evolution*, *Neder-Mead*, *Powell-Brent*, *Best Neighbor*, *Sbplx*, and *Direct*. When the mean 2D projection error is considered, the order changes slightly: *Differential Evolution*, *Powell-Brent*, *Best Neighbor*, *Sbplx*, *Neder-Mead*, and *Direct*. The difference can probably be attributed to the discrepancy between the automatic segmentation used in the alignment process, and the manual segmentation used for computing the error. Nevertheless, the performance of the top ranking algorithms appears to be satisfying. The relatively poor performance of the *Bobyqa* and *Cobyla* algorithms might be an indication that the shape of the energy function is not well represented by their quadratic and linear models, respectively. The *Praxis* method appears to be the least

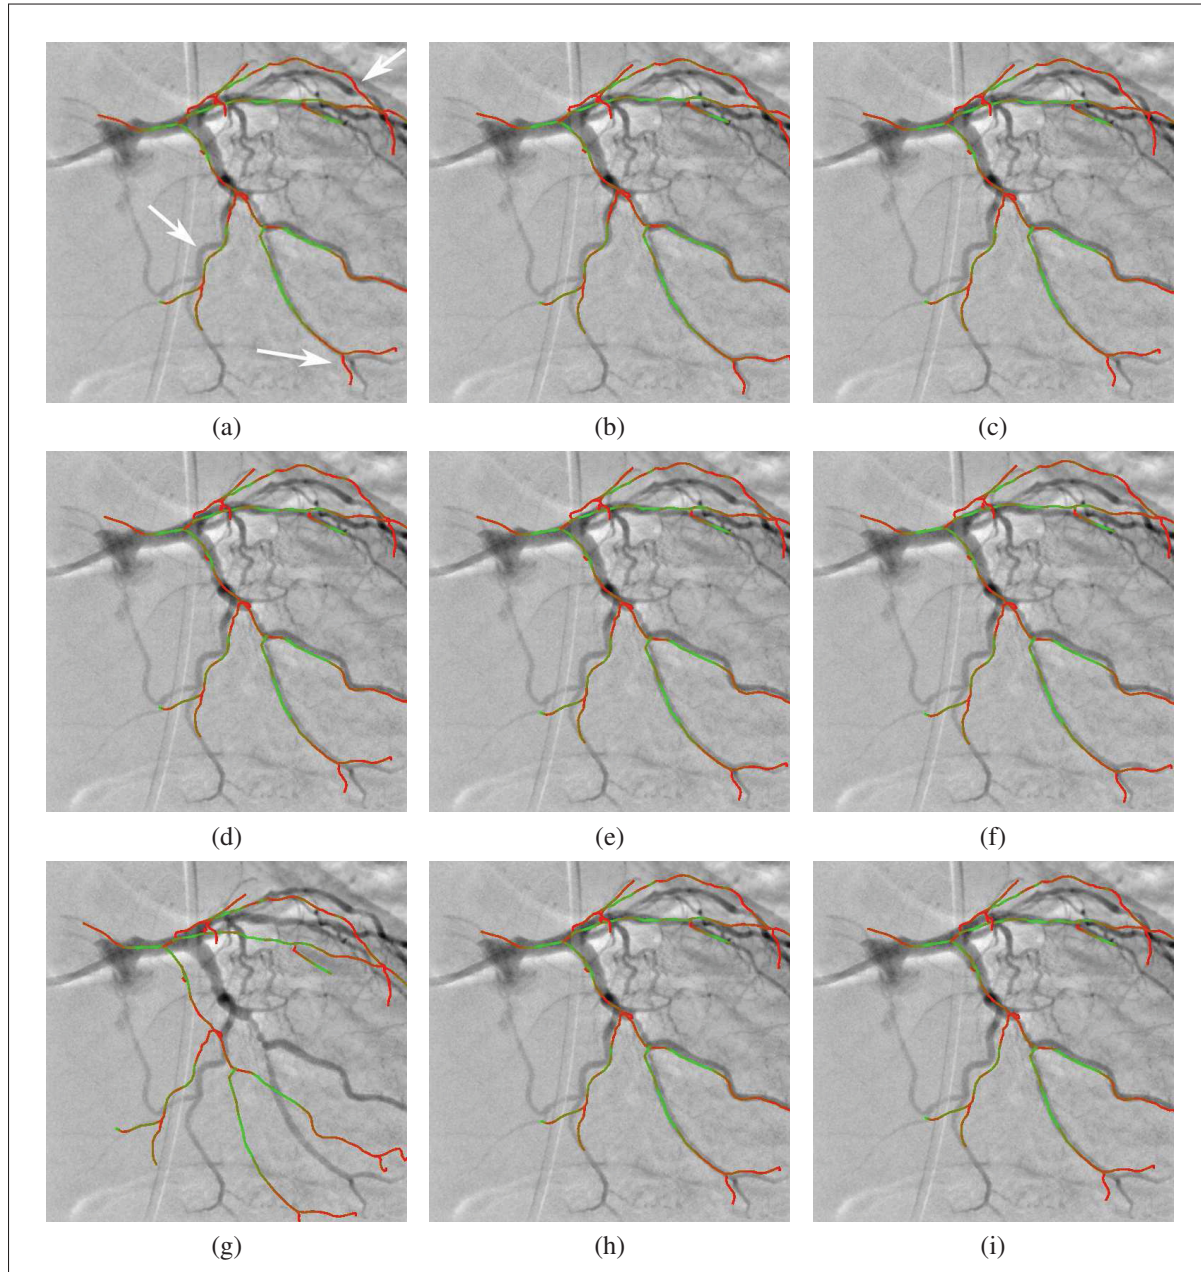


Figure 5.13 Sample output obtained with the affine transformation model using nine different optimizers on the Patient 3 dataset. The values in brackets indicate the residual energy of (5.1). a) *Best Neighbor* [31.86], b) *Differential Evolution* [29.87], c) *Nelder-Mead* [30.70], d) *Direct* [31.69], e) *Bobyqa* [34.90], f) *Cobyla* [36.21], g) *Praxis* [59.68], h) *Sbplx* [31.39], and i) *Powell-Brent* [30.65]. Note that, although only one fluoroscopy plane is shown, the biplane pair was used. The arrows indicate some regions of interest.

appropriate local optimizer for this problem. The box-and-whisker plots in Fig. 5.14 indicate that the performance of the optimizers varies from one dataset to another, even though the pre-

viously discussed observation hold. Thus, the alignment results can probably be boosted by using two or more optimizers in parallel.

The *Differential Evolution* algorithm was found to be the best performer in term of both the residual energy and the mean 2D error. However, the gains in mean 2D error were marginal for a computational time that is about 100 times that of the local algorithms. Still, the median value of the computational time of that optimizer, just under 11 s, might be reasonable for applications that are not time critical. In general, it can be noted that the global optimizers only rarely lead to lower energy or error figures than the best-performing local optimizers. This is an indication that the latter actually finds solutions that are close to the global optimum. It is also worth noting that Fig. 5.14 shows an interesting characteristic of the global optimizers, in that they generally succeed in avoiding the worst solutions. However, because the computational time of the global optimizers is several orders of magnitude larger than that of the local optimizers, we found that they are less appropriate in our setting.

The sample image in Fig. 5.13 shows that there is a good correlation between the residual energy and the perceived visual correspondence. Nevertheless, even with the best alignment with the affine transformation model, a relatively large discrepancy exists between the projected centerline and the 2D fluoroscopies. This can probably be attributed to the presence of non affine deformation in the dataset due to the inaccuracy of the temporal alignment between the modalities, and to the shape change induced by the patient's position change during the various acquisitions steps.

5.5.4 Comparison of the global alignment method with non-rigid registration

In this section, detailed global alignment results obtained using the *Best Neighbor* optimizer are presented and compared with those obtained after non-rigid registration. In all cases, the centerline model was registered with the fluoroscopic frames that are gated in the same cardiac phase used for the CT reconstruction. The results after translational, rigid, affine, and non-rigid registration are shown for the dataset from three patients in Fig. 5.3 and 5.16. The mean projection error of the centerline on the two fluoroscopic planes with respect to 2D manual segmentations has also been calculated (in millimeters). Results are presented in Table 5.3. The computational times needed for the non-rigid registration were in the range [730 – 3 300] ms, with an average of 1 591 ms.

The mean 2D projection errors and their standard deviations decrease, or stay approximately constant, as the complexity of the model increases. However, the relative contribution of the

Table 5.1 RMS residual energy and resulting 2D error, in function of the transformation model.

Optimizer	Translation		Rigid		Affine	
	energy	mm	energy	mm	energy	mm
Best Neighbor	29.98	4.10	22.01	3.56	16.20	2.63
Differential Evolution	29.57	4.05	20.23	3.90	13.69	2.43
Neder-Mead	29.74	4.03	21.24	3.76	15.42	2.78
Direct	29.73	4.07	23.85	3.76	17.38	2.85
Bobyqa	31.20	4.67	22.58	3.55	20.51	3.29
Cobyla	30.12	4.03	22.90	3.43	22.39	3.37
Praxis	46.24	6.75	38.46	5.78	32.90	5.38
Sbplx	30.00	4.15	22.22	3.65	16.69	2.74
Powell-Brent	30.23	4.16	18.05	3.58	15.79	2.60

Note: the values presented above correspond to the RMS results over 195 experiments for each optimizer.

rigid, affine, and non-rigid transformation model varies from one dataset to another. This can be explained by the nature of the deformation presented by each individual case, which is linked to the patient's position, the interval between the CTA and 2D fluoroscopy acquisitions, and the acquisition protocol. Also, patient respiration can cause significant non-rigid heart deformation. The accuracy of the temporal alignment by ECG gating is also important, since the beating of the heart is a significant source of non-rigid deformation. The maximal error also tends to decrease with model complexity, but it sometimes stays stable. In those cases, the performance of the registration algorithm might be limited by an incomplete 2D segmentation, as computed by the automatic segmentation method (Schneider and Sundar, 2010). In fact, in regions where little or no information is available, the non-rigid registration method should alter the centerline as little as possible. This characteristic is desirable in situations where parts of the structure are poorly visible on the X-ray fluoroscopic images, as is sometimes the case with CTO. For an example, see Fig. 5.16–bottom rows: the centerline is well aligned over and under the CTO, but there is little deformation where the contrast is poor.

5.5.5 Global alignment in the multi frame scenario

Another potentially useful scenario, beyond the alignment of a centerline with a biplane pair, is the alignment of a centerline with a temporal sequence of frames from biplane angiography. The performance of the global alignment method in the multi-frames scenario is demonstrated

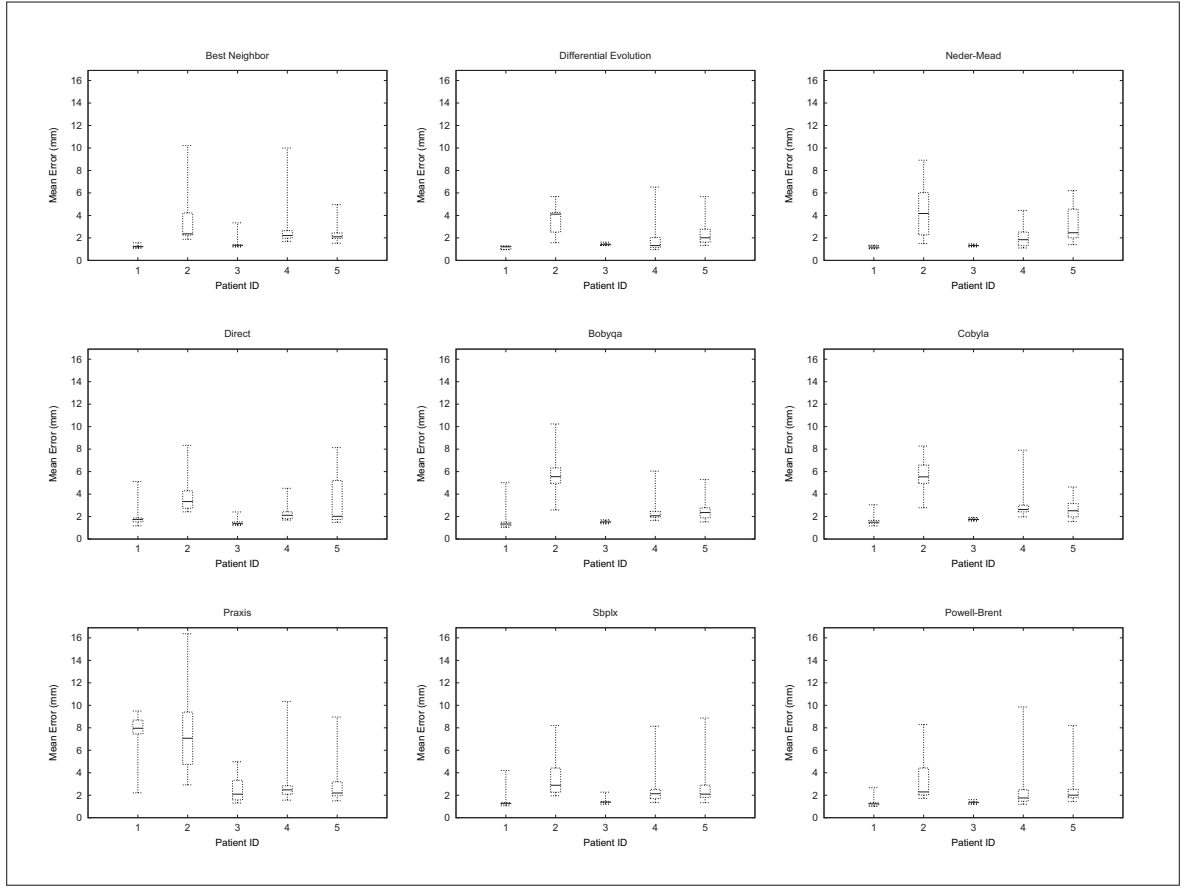


Figure 5.14 Mean 2D error, per optimizer, and per patient, after alignment with the affine transformation model. Each box-and-whisker point correspond to the results of 39 experiments. The boxes represent the first, second, and third quartiles. The whiskers indicate the 5th and the 95th percentiles.

Table 5.2 Total computational time for the two global optimizers.

Optimizer	Min	Q1	Median	Q3	Max
Differential Evolution	6 065	9 177	10 898	11 239	17 079
Direct	81.51	248.4	35 418	84 451	2.604e ⁵

Total computational time, in millisecond, when aligning the 3D centerline using the translation only, rigid, and affine transformation models, successively. Q1 and Q3 correspond to the first and third quartiles respectively.

using a sequence of 11 biplane frames from the Patient 1 LCA dataset. Global alignments were computed using the rigid and affine transformation models. The inter-frame regularization parameters γ in (5.5) varied between 0.0 and 1.0. The *Best Neighbor* optimizer was used in




Figure 5.15 Total computational time, in millisecond, for all local optimizers, when aligning the 3D centerline using the translation only, rigid, and affine transformation models, successively. The boxes represent the first, second, and third quartiles. The whiskers indicate the 5th and the 95th percentiles.

Table 5.3 Mean residual 2D projection error, in mm, calculated after rigid, affine and non-rigid registration for five patients.

	Subject	Translation			Rigid			Affine			Non-Rigid		
		mean	std	max	mean	std	max	mean	std	max	mean	std	max
LCA	1	3.248	3.547	15.661	1.168	1.149	7.031	1.174	1.155	7.031	0.663	0.742	7.031
	2	5.734	4.945	18.621	5.352	4.830	17.959	2.241	2.879	14.440	1.537	1.969	14.440
	3	2.430	2.325	14.398	1.731	1.588	8.700	1.204	1.088	7.886	0.838	0.856	7.886
RCA	4	2.653	2.565	10.205	2.558	2.458	10.037	2.210	2.376	9.190	1.049	1.281	8.243
	5	3.161	2.701	11.701	2.627	2.286	10.259	2.627	2.286	10.259	1.380	1.333	7.251

this case since it was found to be the most effective at minimizing the cost function (5.5). Sample results obtained are shown in Fig. 5.17. Computational times were 39 s, 89 s, and 91 s, for the experiments presented in the top, middle, and bottom row respectively. As anticipated, it was found that using the affine transformation model led to better qualitative results than using the rigid one. Using regularization ($\gamma = 0.1$) also improved the results when using the affine transformation model. With large regularization values, $\gamma > 1.0$, the stiffness of the model




Figure 5.16 From left to right: input image, rigid alignment, affine alignment, and non-rigid registration. Both biplane images are shown in all cases. Top and bottom rows correspond to Patients 1 and 2 in Table 5.3, respectively. The arrows indicate some regions of interest.

increases, and the optimizers cannot find a local minima that is far from the initial position. Preliminary investigation suggests that domain specific optimizers might achieve better results, e.g. by allowing the parameters to change in a coordinated manner. Nevertheless, we found that the current setup was appropriate for tracking the alignment transformations of the LCA during most of the cardiac cycle. Experiments on the RCA datasets did not lead to convincing

results. In fact, even though the global alignment procedure works, the amount of non-rigid deformation sustained by the RCA during the cardiac cycle renders the affine model ineffective except when the CTA and X-ray acquisition are in close temporal alignment.




Figure 5.17 Registration over a sequence of frames. *Top row*) rigid transformation model [$\gamma = 0.1$, $E_{\text{Multi}}^* = 120.284$]; *middle row*) affine transformation model [$\gamma = 0$, $E_{\text{Multi}}^* = 119.356$]; and *bottom row*) affine transformation model + regularization [$\gamma = 0.1$, $E_{\text{Multi}}^* = 116.905$]. The arrows indicate some regions of interest, E_{Multi}^* corresponds to the first righthand term of (5.5).

5.5.6 Semiautomatic tracking of the right coronary artery.

In this experiment, a semiautomatic procedure based on the proposed non-rigid registration method has been used to track the RCA. Starting from the gated frame, global alignment and non-rigid registration are performed. This deformed centerline model is then used by the operator as the initial model and position for the next pair of fluoroscopic frames. This process is repeated for all frames over one cardiac cycle, as presented in Fig. 5.18 and in the attached video. As can be seen, this procedure permits successful tracking of the RCA during one cardiac cycle. Because this experiment has been conducted before the simulation study, a slightly different set of parameters was used: $\mu = 0.05$, $\nu = 5.0$, and $\lambda = 0.1$. The required amount of time is about 15 minutes, which is reasonable in an offline setting.

5.6 Multimedia Material

Three videos illustrating the method are available at www.synchromedia.ca/reg2D3D. They present: 1) an iterative update of the non-rigid registration; 2) the full alignment and registration progress; and 3) the tracking of an RCA artery using non-rigid registration.




Figure 5.18 Tracking the RCA. From left to right: at the gated frame, and at +4, +6, +8, and +10 frames. The initial curve is green/red, and the registered curve is blue.

5.7 Discussion and conclusion

A new 2D/3D registration method has been proposed, and applied to the problem of registering a 3D centerline model of the coronary arteries with a pair of fluoroscopic images. The methodology is divided into two main parts: 1) global alignment, and 2) non-rigid registration.

In the first part, an energy depending on a global transformation model (translation-only, rigid, or affine) is defined. Nine general purpose optimization algorithms have been used to minimize this energy, which results in an estimation of the 2D/3D alignment transformation parameters. Based on the experiments on clinical data, it appears that the following local optimizers give good results in a short time (median < 105 ms, max < 950 ms): *Neder-Mead*, *Powell-Brent*, *Best Neighbor*, and *Sbplx*. Considering also the simulation results, the *Neder-Mead* algorithm was the best overall performer. The two global optimizers only rarely led to major improvements in the result and required a computational time orders of magnitude higher. Disregarding computational time, the *Differential Evolution* algorithm generally returned the best solution. When using a local optimizer, the alignment time was consistently under 1 s, which makes the method suitable for use during an intervention. The advantage of using an affine transformation instead of a rigid transformation is dependent on the nature of the dataset and can be significant in some cases. Overall, it was found that the global alignment procedure is appropriate for use on both LCA and RCA datasets, when the 3D and 2D modalities are temporally aligned using ECG gating. The experiment on the dataset from a CTO patient demonstrates the benefit of the



proposed method when applied to similar clinical cases. The proposed method can also help with intervention guidance by augmenting the 2D imagery with the 3D geometry segmented from a CTA acquisition, thereby greatly reducing the ambiguities inherent in the interpretation of the 2D images.

As for non-rigid registration, the proposed method uses a robust reconstruction strategy to compute the forces used to deform the 3D model. Regularization terms limit the total displacement of the segments, ensure that the displacements are smooth, and help preserve the relative orientation of the branches, making the non-rigid registration results plausible. The regularizers included in the deformation model enable the elegant management of regions with missing data (e.g. Patient 2 in Fig. 5.16) by deforming the centerline a minimum amount to ensure coherence with the rest of the structure. No excessive deformation or straightening will occur. Experiments on five different patients were presented with promising results. The total computational time, which was generally below 3 s, is acceptable for interactive applications. Nonetheless, better numerical algorithms with faster computational time will be researched. The proposed non-rigid registration method will make the centerline snap to adjacent 2D structures when started from an appropriate initial point, making interpretation of the 2D images easier, especially for difficult low contrast CTO cases. This also makes it possible to present an updated 3D model alongside the operational images to provide an improved perception of the 3D space.

Finally, experiments using multiple X-ray biplane angiography frames have also been presented. It was found that the multiframe global alignment method works well on LCA datasets, and the proposed interframe regularizer leads to improved results. A non-rigid semiautomatic tracking procedure has been devised to handle cases with more non-rigid deformation, such as with non-temporally aligned RCA datasets, and applied to one such dataset. The semi-automatic method required approximately 15 min of interaction for a sequence covering one heart beat, which seems reasonable for practical offline applications. Future work will focus on automating this non-rigid method for tracking coronary arteries.

Acknowledgment

The C++ code for DRR generation has been generously provided by Julien Couet and Luc Duong of École de technologie supérieure.

CHAPTER 6

DISCUSSIONS

The general objective of this thesis was to define a framework for 3D segmentation in medical imaging scans and alignment of the segmented structure model with interventional fluoroscopy for intervention guidance. In this context, our proposed general methodology covers three important aspects of this intervention guidance framework, all of which are complementary. Chapter 3 presented our work on the definition of a new local-linear level-set method for the segmentation of large structures with spatially varying intensity. Chapter 4 introduced a new local structure model and defined an image filter for the detection of small curvilinear structures. Finally, 2D/3D registration of centerline models with live fluoroscopies was the subject of chapter 5. All this work has been published as independent journal articles, in order to disseminate them as widely as possible. However, it is interesting to note how well those individual studies fit together in the proposed general framework. Below, we discuss the place occupied by each one in that framework, and highlight the strengths and weaknesses of the proposed techniques.

6.1 A level-set method using local-linear region models for the segmentation of structures with spatially varying intensity

In chapter 3, we defined a level-set method that uses a local-linear region model. With this new region-based formulation, not only is the fitness of voxel intensities tested, but the fitness of the spatial derivatives of the voxel intensities is tested as well. Our method has the potential model image regions with local gradients, something that was not possible with previously defined models — see figure 3.2 on p.61. Furthermore, this region model has been justified probabilistically, and is shown to be appropriate when low-frequency image gradients play a significant role in the definition of the various image regions. To illustrate value of the model, picture a long hall, where only the end closest to us is lit: the brightness of light reflected on the wall slowly diminishes as we look further down the hall. In this picture, it would generally make more sense to segment the walls, the floor, and the ceiling rather than to partition the image in bright and dark regions. Our model would allow to correctly segment this image, and it allows to segment tissues in the same way. Since the method makes use of the local image gradients during the evolution of the level-set, special attention must be paid to the presence of image noise. For example, the local gradient could be estimated using Sobel's operator or a denoising filter might be used on the image before starting the segmentation process. However,

embedded into our definition of the method is an edge indicator function, which tends to 0 as the image gradient increases and limits the potential damage that could be caused by spurious noise. This method has been defined in an article published by Elsevier in the journal *Magnetic Resonance Imaging*, and is the main contribution of this article.

The derivation of the Euler-Lagrange equation of the proposed energy functional for segmentation with local-linear region models results in a relatively complex level-set evolution equation. Not only does this complicate the implementation of the method, but integral terms significantly slow down the computational process. Consequently, we have derived a fast approximate flow for the minimization of the proposed level-set functional, which approximatively halves the computational time. This is another important contribution of this article.

Because of its proprieties, the benefits of our local-linear method appear more clearly in the segmentation of tissues in MRI. This is why we tackled the challenging problem of segmenting brain tissues in MRI in our research.

6.2 Brain tissues segmentation in MRI

MR images are subject to various electromagnetic perturbations, which appears as low frequency differences in imaging intensity occurring on the target domain, and causing what it is known as a *bias field*. Even though it is possible to visualize and even quantify the bias field by imaging an object with constant MR properties, the complex relationship between the bias field and the imaged object makes it difficult to compensate for its effect retrospectively. To further complicate the issue, the literature indicates that the bias field can affect each class of brain tissues differently. In fact, this significant source of intra-scan intensity nonuniformity (INU) has made automatic brain tissue segmentation a challenging task.

Since our new local-linear level-set method uses a different model for each region, it appears particularly well suited for the segmentation of brain tissues with a significant amount of INU. In the experiments presented, this method was found to be very successful at segmenting the complex 1.5 Tesla scans from the IBSR database, with Tanimoto indexes ranging from 0.61 to 0.79 for the classification of white matter, and from 0.72 to 0.84 for gray matter. As indicated in chapter 3, these results are significantly higher than other published results on the same database, which means that the definition of an automatic method for brain tissue segmentation in MRI scans with significant INU levels is another important contribution of this research. It is also worth noting that, to the best of our knowledge, this is the first time a region-based level-set model has been used to perform the segmentation of real world MRI brain scans

with convincing results. Finally, we plan to assess the performance of the method for the segmentation of other biological structures in the future. However, it might be interesting to note that the local-linear level-set method has already been proved useful to segment other type of shapes in a widely different context (Rivest-Hénault *et al.*, 2011).

6.3 Accurate local structure modeling and vessel detection using *structure balls*

A major problem in the segmentation of vascular structures in CTA or MRA scans of the chest is segmenting the very fine distal vessel segments, which can be less than one or two voxels in diameter. Global techniques, such as level-set or graph-based methods, are generally ineffective for two main reasons: 1) their regulation (smoothing) term penalizes regions with a thin profile, and 2) the partial volume effect in the scan distorts the vessel intensity observed, which confuses the algorithm. Two principal types of approach can be found in the literature to segment these difficult structures: tracking approaches, and local detection approaches. Tracking approaches are generally more accurate, but they usually require a user to indicate starting points and to indicate the approximate vessel diameter. In contrast, vessel detection approaches are more general in that they are able to automatically detect most vessels of a given volume in one step. Unfortunately, most existing vessel detection methods return poor results around vessel bifurcations, which can be problematic if the vessel detection signal is to be used by a subsequent process.

To address this issue, we have defined a new local vessel indicator based on the analysis of a second order difference operator applied in many directions, a geometric construction that we call a *structure ball*. This construction is entirely novel in the context of vessel detection, and its definition in chapter 4 is the main contribution of this work. In addition, many supporting techniques have been introduced, and which also represent significant contributions. A visually appealing technique for the visualization of the image structure makes it possible to directly inspect the *structure ball*'s behavior for any given image region, which in turn provides insight into the filter mechanisms involved and could certainly help in devising subsequent tracking or segmentation methods. In addition, a compact representation of the *structure ball* using spherical harmonics (SH) is provided, which allows the data to be stored in a memory efficient way. This SH representation is also a natural means of interpolating angular information. Various shape descriptors adapted to the detection of blood vessels are proposed and computed on representative vessel and non vessel patterns. Analysis of the response of those shape descriptors has allowed us to define a vessel detection function that is specific to the vessel pattern, as well as being robust to the presence of bifurcations along the vessels.

In a paper that has been submitted for publication in *IEEE Transactions on Image Processing*, the resulting vessel detection filter has been compared with state-of-the-art methods on both synthetic and real medical images. Based on this comparison, we find that the benefits of this new approach are the following: better sensitivity and specificity characteristics, improved separation of adjacent structures, and the ability to represent vessel junctions. Since similar vessel detection techniques have been used as a building block in many segmentation, tracking, and 2D/3D registration methods, the impact of this improved method is potentially significant, as better vessel detection may enhance the final results.

6.4 2D/3D registration of centerline models with X-ray angiography using global transformation models

Once the structures of interest have been segmented in 3D, and a 3D centerline model has been created from this segmentation, it is possible to enhance navigation guidance by accurately aligning the centerline model with the structures visible on the interventional fluoroscopies. Unfortunately, the appropriate 2D/3D registration methods that could be found in the literatures are too slow, with computational times in the order of few minutes, or too inaccurate to be used during an intervention. In research carried out in collaboration with Siemens Corporate Research (SCR), we developed a novel method for non-rigid 2D/3D registration of coronary artery models with live fluoroscopy that is accurate and fast enough to enhance cardiac intervention guidance. This method is composed of two steps: 1) global alignment using a rigid or an affine transformation method, and 2) non-rigid registration. This work has been formally evaluated, and our findings are soon to appear in *IEEE Transactions on Medical Imaging* (in press). In addition, the method has been implemented in C++ within the Siemens Inspace imaging platform. Because the optimization problems involved in the first global alignment step and in the non-rigid registration step differ significantly, we have broken down our discussion into two parts. The first step follows below, and the second step is the subject of the following section.

Most of the literature on 2D/3D registration was developed around the registration of rigid anatomical structures, such as bones, which probably explains why rigid transformation models are most often used. When deformable structures are considered, such simple transformation models are limited, and can be incapable of generating a convincing registration. For example, the coronary arteries undergo significant shape change during the cardiac cycle, and any small temporal miss alignment between the CT acquisition and the X-ray angiography would impact the registration results. In addition, CT is acquired under a breath hold, while fluoroscopy is

generally acquired in a free breathing state, a difference that is known to increase the shape dissimilarities observed. This makes it very important to use a transformation model that is more flexible than the rigid transformation model. In our research, we began by using transformation models of increasing complexity to compute the initial alignment of the centerline model with biplane fluoroscopy: translation, rigid, and affine. We found that the contribution of the affine model to the final alignment varies from one dataset to another, but can be significant in some cases. This is probably related to the temporal alignment of the modalities and to the breathing style of the patient.

At the processing level, when compared to existing rigid methods, adding the affine transform is significant since the added complexity generally involves much more effort from the optimizer, and slows down the process. Our method remedies this problem by using a cost function that can be computed extremely rapidly. In fact, we use the distance transform of the coronary arteries segmented in 2D on the X-ray fluoroscopies. These distance transforms need only to be computed once per image, and can be precomputed. This makes evaluating the cost function simply a matter of picking values in a look-up table.

It has also been suggested in the literature that the optimizer chosen to minimize the alignment cost function might have a significant impact on the final result. Surprisingly, most authors have only reported the results they have obtained with one or two optimizers. This means that our evaluation of the performance of nine general purpose optimizers for estimating the alignment parameters is also an important contribution of this research. This evaluation revealed significant differences between the optimizers. The Nelder-Mead algorithm (also known as the downhill simplex) was the best overall performer, with the lowest computational cost and results that were close to the best. For offline applications, our experiments suggest that a global method, such as the differential evolution, could generate slightly better results and be less sensitive to the initial position, at the expense of a computational time several orders higher.

6.5 Non-rigid 2D/3D registration of centerline models with X-ray angiography

The affine alignment method discussed above is, in many cases, an improvement over rigid methods. However, there are still many situations where the computed affine transformation is not sufficient to convincingly overlay the 3D centerline model on the 2D fluoroscopies. We proposed a new efficient non-rigid 2D/3D registration method that can be used to compensate for the remaining shape discrepancies after the initial affine alignment. Compared with the few other non-rigid techniques published in the literature, this method is unique, in that it is not

tied to the registration of any one type of structure and it has a short runtime. In this respect, our principal contribution is certainly the definition of a non-rigid registration method as a simultaneous matching, reconstruction, and registration problem, that can, on modern general purpose hardware, be solved fast enough to be used intraoperatively during CTO procedures. In our experiments, typical runtimes were under 3 s. In tests conducted using both synthetic and real world images, we found that the non-rigid method is able to compensate for relatively high differences in position, up to 6mm in some cases, which is significant, since this is many times the diameter of imaged coronary arteries. Also, experiments have shown that the method is robust to the presence of Gaussian noise and segmentation errors.

6.6 Improved surgical guidance using 2D/3D registration

In a typical CTO scenario involving percutaneous coronary intervention, the practitioner will periodically turn on the X-ray beam and inject contrast agent in order to evaluate the position of the instrument with respect to the target structures. Because both the X-rays and the contrast agent are harmful in large doses, the medical team tries to minimize patient exposure to these agents. In this context, the non-rigid 2D/3D registration method proposed in this thesis can be used to compute the transformation necessary to overlay a patient-specific model on the interventional imagery, thereby improving surgical guidance by continuously presenting an accurate representation of the patient structures in the same reference frame as the surgical instruments. For improved accuracy, the registration method makes it possible to periodically update the transformation of the 3D centerline model, for example after every fluoroscopy acquisition. This scenario has been presented to a radiologist at CHU Ste-Justine Hospital, and has been tested at a Siemens clinical test site located in The Netherlands. In both cases, we received enthusiastic encouragement.

The periodic updating of the centerline model can be quite fast, and the speed of the algorithm can be considered as *interactive*. On a laptop with an Intel i7 Q820 processor running at 1.73 GHz, the total registration time was typically less than 3 seconds plus 280 milliseconds per image for the precomputed 2D segmentation and distance transform. If we consider that fluoroscopies for surgical guidance are typically acquired at a rate of 15 frames per second (fps), a soft real time algorithm would require a processing time of less than 66.7 milliseconds for each frame¹. In this respect, the proposed algorithm is not real time, but, a closer look shows that it is not far from this target. Indeed, the 2D segmentation algorithm used to pre-process the fluoroscopies is already parallelized, and running this algorithm on a dedicated workstation

¹This would introduce a lag of up to 66.7 ms into the visualization pipeline, which is probably acceptable.

with better, more recent, commodity processors (e.g. $4 \times$ Intel Xeon E3-1290 v2 at 3.7 GHz) could easily reduce the computational time by a factor of 8, at a meager estimated cost of 35 milliseconds. Since some image filtering is involved with this algorithm, it is also likely that a fast GPU implementation will also help. For the alignment with the global transformation models, the median runtimes are already quite low on the i7 laptop with a single-threaded C++ implementation. The translational transform takes about 5 milliseconds to compute, the rigid one 10 milliseconds, and the affine one 30 milliseconds, for a total of 45 milliseconds for the full sequence. In this case, using a dedicated workstation and an efficient multi-threaded implementation would allow to create a margin of safety for the worst cases.

In terms of the non-rigid step, reducing the computational time from about 3 seconds to a small fraction of a second would require much more effort. However, the question is: is real time 2D/3D non-rigid registration really needed? In a clinical setting, the most obvious application for a continuously morphing model is to provide some sort of visual feedback that temporally interpolates the position of the structure between the periodic fluoroscopy acquisitions. In this case, the strategy could be to create a 3D+t model before the procedure, and define a robust registration and interpolation method to synchronize this 3D+t model with the acquired fluoroscopies during the operation. A semi-automatic technique to create such a 3D+t model has been demonstrated in chapter 5. However, work is still needed to define the registration and interpolation method required to convincingly overlay the dynamic model on the interventional images. Integration of such a technique with a non-invasive instrument tracking device has the potential to define a new guidance standard. Such a system could potentially decrease interventional risks, and reduce patient exposure to X-rays and contrast agent.

6.7 Application for MAPCAs procedure

A similar navigation guidance framework has been proposed in collaboration with partners from CHU Ste-Justine Hospital. First, we proposed a different, but related level-set based method for the segmentation of the aorta and attached major aorto-collateral arteries (MAPCAs). The main feature of the method is that it benefits from user-provided seed points to pre compute a spatially varying region intensity model. When compared with the level-set method with local-linear region model presented in chapter 3, the main advantage of this new method is that it needs less computational time, and only the structure of interest is modeled. The latter difference helps prevent a very uneven background from interfering with the segmentation of the structures in thoracic CTA. This work has been implemented in C++ using Siemens's

Xip platform and ITK and was presented at SPIE Medical Imaging 2010. For the sake of completeness, this article is also presented in Appendix II.

The 3D segmented model has also been used in a navigation guidance framework in pediatric cardiology. In this respect, a 3D model-to-image registration strategy for registering this model with interventional fluoroscopies has been proposed by one of my colleagues, Julien Couet, and an article on this method has been recently presented at SPIE Medical Imaging 2012 (I am one of the coauthors of this article). In addition, the method has been presented to a group of clinicians and researchers at CHU Ste-Justine Hospital and was received enthusiastically. This image-based registration method is more computationally intensive than the model-to-model method presented in chapter 5, and its transformation model is less flexible. However, since the MAPCAs are mostly affected by respiration, a motion that causes slower and less dramatic shape changes than the beating of the heart, a simple rigid transformation model gave satisfactory results for this problem, and contributed to keeping the computational time at an acceptable level. This second application demonstrates the versatility of similar enhanced image-guided navigation frameworks.

GENERAL CONCLUSION

This thesis presented new methods for the segmentation and 2D/3D registration of medical images. The common denominator of the methods presented is the key idea of localization, and the use of differential geometry. By taking into account the information available in the neighborhood of each 2D or 3D image pixel, it was possible to define methods that produce better results on difficult real world images. We believe that together these methods constitute a step forward in the establishment of diagnostic, surgical planning and guidance frameworks. Both the proposed local-linear level-set segmentation method and the proposed vessel detection filter are robust to local changes in the intensity of the voxels corresponding to the structures of interest, a problem associated with the limitations of the imaging equipment, and also with the propagation of the contrast agent in blood vessels. Further more, our newly introduced vessel detection filter makes it possible to automatically extract curvilinear structures, even in low contrast regions of 3D CTA and MRA scans. Last, but not least, the proposed 2D/3D non-rigid registration method enable the 3D information of a 3D CTA scan to be overlaid on 2D interventional images, which partly remedies to the limitations of X-ray fluoroscopy, and enhances surgical intervention guidance.

At the clinical level, the contributions of this thesis allow us to envision many possibilities. The automatic brain tissue segmentation algorithm based on level-set and local-linear region models has the potential to reduce human interaction in morphological studies, thereby increasing their reproducibility. Closer to our original application context, the proposed level-set methods have the potential to greatly facilitate the segmentation of the vascular structures in difficult pediatric CTA scans. The proposed vessel detection filter is a significant improvement over existing techniques, and is suitable for the segmentation of smaller vessels. A hybrid formulation, using local region model and vessel detection, could certainly be devised to bring together the benefits of these two complementary methods, allowing practitioners to produce accurate, and repeatable segmentations, and do so more quickly.

The non-rigid 2D/3D method for the registration of coronary arteries with live fluoroscopy offers a practical solution for the guidance of difficult laparoscopic cardiac interventions. Based on the results presented, it is possible to conceive of similar 2D/3D registration methods that could serve for guidance in other procedures, as has been discussed in the previous chapter. In addition, the results obtained pave the way to the creation of future 3D+t patient specific models that could be continuously updated during an intervention. It is our fervent hope that

such technologies will reduce the risk for the patient and increase the success rate of difficult interventions.

Summary of contributions

In this section, we briefly highlight the major contributions of this thesis.

1. A new local-linear level-set model has been introduced and is, to the best of our knowledge, the first level-set method that is capable of segmenting regions with significant and potentially local gradients. The performance of this new model has been assessed on two very different problems: automatic brain tissue segmentation and document image binarization. For both applications, it was the first time that a level-set method was used and formally evaluated for this task with convincing results on publicly available databases.
2. The *structure ball*, a novel local structure model for 3D graylevel images, has been defined and used in the context of curvilinear structure detection. The resulting image filter is the first of its kind that allies strong discrimination characteristics with the possibility to take into account more than one local structure directions. In the presented experiments, this new curvilinear structure detection filter is an improvement over the well known Frangi's filter (Frangi *et al.*, 1998) in every aspects.
3. A complete and practical 2D/3D non-rigid registration method for the alignment of 3D patient specific centerline models with biplane angiography has been introduced. With a typical runtime of less than 3 seconds, this method is, to the best of our knowledge, the first 2D/3D non-rigid registration method that is fast enough to be used interactively during cardiac interventions. It has been implemented in Siemens Inspace medical imaging platform, and tested at an experimental clinical site located in The Netherlands with promising results.

Articles in peer reviewed journals

1. D. Rivest-Hénault and M. Cheriet, "Vessel detection filter via structure-ball analysis". (Submitted) *IEEE Transactions on Image Processing*.
2. D. Rivest-Hénault, H. Sundar and M. Cheriet, "Non-Rigid 2D/3D Registration of Coronary Artery Models with Live Fluoroscopy for Guidance of Cardiac Interventions". In *IEEE Transactions on Medical Imaging*, Volume 31, Number 8, August 2012, Pages 1557–1572. <http://dx.doi.org/10.1109/TMI.2012.2195009>.

3. D. Rivest-Hénault and M. Cheriet, “Unsupervised MRI segmentation of brain tissues using a local linear model and level set”. In *Elsevier Magnetic resonance imaging*, Volume 29, Number 2, February 2011, Pages 243–259. <http://dx.doi.org/10.1016/j.mri.2010.08.007>.
4. D. Rivest-Hénault, Reza Farrahi Moghaddam and M. Cheriet, “A local linear level set method for the binarization of degraded historical document images”. In *International Journal on Document Analysis and Recognition*, Volume 15, Number 2, 2012, Pages 101–124. <http://dx.doi.org/10.1007/s10032-011-0157-5>.

Articles in conference proceedings with a reading committee

1. J. Couet, D. Rivest-Hénault, J. Miró, C. Lapierre, L. Duong and M. Cheriet, “Intensity-based 3D/2D registration for percutaneous intervention of major aorto-pulmonary collateral arteries”. In *SPIE Medical Imaging 2012*. (San Diego, February 4–9 2012). SPIE Publishing.
2. D. Rivest-Hénault, C. Lapierre, S. Deschênes and M. Cheriet, “Length increasing active contour for the segmentation of small blood vessels.” In *20th International Conference on Pattern Recognition (ICPR 2010)*. (Istanbul, August 23–26 2010). IEEE Computer Society.
3. D. Rivest-Hénault, L. Duong, C. Lapierre, S. Deschênes and M. Cheriet, “Semi-automatic segmentation of major aorto-pulmonary collateral arteries (MAPCAs) for image guided procedures.” In *SPIE Medical Imaging 2010*. (San Diego, February 13–18 2012). SPIE Publishing.
4. R. Farrahi Moghaddam, D. Rivest-Hénault, I. Bar-Yosef and M. Cheriet, “A unified framework for the restoration of bleed-through problem in double-sided document images based on the level set approach.” In *10th International Conference on Document Analysis and Recognition (ICDAR 2009)*. (Barcelona, July 26–29 2009), pp. 441–445. IEEE Computer Society.
5. R. Farrahi Moghaddam, D. Rivest-Hénault and M. Cheriet, “Restoration of highly degraded characters with minimum prior information using a level set approach enhanced with multi-level classifiers.” In *10th International Conference on Document Analysis and Recognition (ICDAR 2009)*. (Barcelona, July 26–29 2009), pp. 828–832. IEEE Computer Society.

6. D. Rivest-Hénault and M. Cheriet, “Image segmentation using level set and local linear approximations.” In *International Conference on Image Analysis and Recognition (ICIAR 2007)*. (Montreal, August 22–24), pp. 234–245. Springer Lecture Notes in Computer Science (LNCS).

Internship at Siemens corporate research

I did a seven months internship at Siemens Corporate Research (SCR). During that time I developed the method presented in chapter 5. In addition to this work, I implemented and improved a novel visualization method for 3D medical CT scan in Siemens Inspace platform using C++ and OpenGL. The method makes a 2D rendering of a 3D medical scan by slicing it in a way that eases the visualization of complex anatomical structures, such as the coronary arteries.

Open source vessel detection filter

A C++/ITK code implementing the vessel detection filter described in (Rivest-Hénault and Cheriet, 2012) is now freely available on the web.

DIBCO’2009 Contest

Third best performing document image binarization algorithm out of 43 algorithms submitted by 35 international teams at the DIBCO’09 document image binarization contest held under the framework of ICDAR 2009.

Paper reviewing

I reviewed scientific papers for the following journals and conferences:

- IEEE Transactions on Medical Imaging
- International Conference on Information science, signal processing and their applications, ISSPA2012
- International Conference on Document Analysis and Recognition. ICDAR 2009 and 2011.

APPENDIX I

EULER-LAGRANGE EQUATION FOR LOCAL LINEAR STATISTICS

In this section, we derive the Euler-Lagrange equation for functional (3.10) and we describe the approximation that has been used in this paper. Empirical evidences of the validity of the approximation are also presented.

Let M_i be a membership function defined as $M_+ = H(\phi)$ and $M_- = (1 - H(\phi))$. The functional that we want to minimize is given by (3.10) and can be re-written as follows:

$$F(a_i, b_i, c_i, \phi) = \sum_{i \in +, -} \int M_i(\mathbf{x}) F_i(\mathbf{x}, a_i, b_i, c_i, \phi) d\mathbf{x} \\ + \nu \int |\nabla H(\phi)| d\mathbf{x},$$

with

$$F_i(\mathbf{x}, a_i, b_i, c_i, \phi) = \alpha(a_i(\mathbf{x}) + b_i(\mathbf{x})x + c_i(\mathbf{x})y - u_0(\mathbf{x}))^2 \\ + \mu g(u_0(\mathbf{x})) \left[\left(b_i(\mathbf{x}) - \frac{\partial u_0}{\partial x}(\mathbf{x}) \right)^2 + \left(c_i(\mathbf{x}) - \frac{\partial u_0}{\partial y}(\mathbf{x}) \right)^2 \right].$$

Neglecting the viscosity term, since it is independent of the region statistics, we compute the Euler-Lagrange equation

$$\frac{\partial}{\partial \phi} \left[\sum_{i \in \{+, -\}} M_i(\mathbf{x}) F_i(\mathbf{x}, a_i, b_i, c_i, \phi) \right] = 0,$$

which terms expand to

$$\frac{\partial}{\partial \phi} M_i(\mathbf{x}) F_i(\mathbf{x}, a_i, b_i, c_i, \phi) = \tag{I.1} \\ \underbrace{\frac{\partial M_i(\mathbf{x})}{\partial \phi} F_i(\mathbf{x}, a_i, b_i, c_i, \phi)}_{\text{Usual term}} + \underbrace{M_i(\mathbf{x}) \frac{\partial F_i(\mathbf{x}, a_i, b_i, c_i, \phi)}{\partial \phi}}_{\text{Supplementary term}}.$$

As it will be discussed below, the *usual* term is the only one used in this paper for the minimization of our functional. The *supplementary* term contain finer information and have been

neglected in our implementation. For the *usual* term, it is easy to see that

$$\frac{\partial M_+(\phi(\mathbf{x}_2))}{\partial \phi(\mathbf{x}_2)} = +\delta(\phi) \quad \text{and} \quad \frac{\partial M_-(\phi(\mathbf{x}_2))}{\partial \phi(\mathbf{x}_2)} = -\delta(\phi),$$

where $\delta(\phi) = \frac{\partial H(\phi)}{\partial \phi}$, so that this term expands to

$$\begin{aligned} \delta(\phi) & \left(\alpha(a_- + b_-x + c_-y - u_0)^2 \right. \\ & \quad + g \cdot \mu \left[\left(b_- - \frac{\partial u_0}{\partial x}\right)^2 + \left(c_- - \frac{\partial u_0}{\partial y}\right)^2 \right] \\ & \quad - \alpha(a_+ + b_+x + c_+y - u_0)^2 \\ & \quad \left. - g \cdot \mu \left[\left(b_+ - \frac{\partial u_0}{\partial x}\right)^2 + \left(c_+ - \frac{\partial u_0}{\partial y}\right)^2 \right] \right), \end{aligned}$$

which, omitting the viscosity term, corresponds exactly to the result presented in the level set equation (3.11).

We are then left with the second supplementary right hand term. With the abbreviation

$$\begin{aligned} F_{i1}(a_i, b_i, c_i, \phi) &= (a_i(\mathbf{x}) + b_i(\mathbf{x})x + c_i(\mathbf{x})y - u_0(\mathbf{x}))^2, \\ F_{i2}(a_i, b_i, c_i, \phi) &= \left(b_i(\mathbf{x}) - \frac{\partial u_0}{\partial x}(\mathbf{x}) \right)^2, \\ F_{i3}(a_i, b_i, c_i, \phi) &= \left(c_i(\mathbf{x}) - \frac{\partial u_0}{\partial y}(\mathbf{x}) \right)^2, \end{aligned}$$

we compute

$$\frac{\partial}{\partial \phi} F_i(a_i, b_i, c_i, \phi) = \alpha \frac{\partial}{\partial \phi} F_{i1} + \mu g(u_0) \frac{\partial}{\partial \phi} F_{i2} + \mu g(u_0) \frac{\partial}{\partial \phi} F_{i3},$$

which leads to

$$\begin{aligned}
\frac{\partial}{\partial \phi} F_{i1} &= \frac{\partial}{\partial \phi} (a_i(\mathbf{x}) + b_i(\mathbf{x})x + c_i(\mathbf{x})y - u_0(\mathbf{x}))^2 \\
&= 2 \cdot (a_i(\mathbf{x}) + b_i(\mathbf{x})x + c_i(\mathbf{x})y - u_0(\mathbf{x})) \cdot \frac{\partial}{\partial \phi} (a_i(\mathbf{x}) + b_i(\mathbf{x})x + c_i(\mathbf{x})y), \\
\frac{\partial}{\partial \phi} F_{i2} &= \frac{\partial}{\partial \phi} \left(b_i(\mathbf{x}) - \frac{\partial u_0}{\partial x}(\mathbf{x}) \right)^2 \\
&= 2 \cdot \left(b_i(\mathbf{x}) - \frac{\partial u_0}{\partial x}(\mathbf{x}) \right) \cdot \frac{\partial b_i(\mathbf{x})}{\partial \phi}, \\
\frac{\partial}{\partial \phi} F_{i3} &= \frac{\partial}{\partial \phi} \left(c_i(\mathbf{x}) - \frac{\partial u_0}{\partial y}(\mathbf{x}) \right)^2 \\
&= 2 \cdot \left(c_i(\mathbf{x}) - \frac{\partial u_0}{\partial y}(\mathbf{x}) \right) \cdot \frac{\partial c_i(\mathbf{x})}{\partial \phi}.
\end{aligned}$$

It is then needed to compute the values for $\frac{\partial a_i(\mathbf{x})}{\partial \phi}$, $\frac{\partial b_i(\mathbf{x})}{\partial \phi}$, and $\frac{\partial c_i(\mathbf{x})}{\partial \phi}$. This can be done by solving

$$A \frac{\partial \theta}{\partial \phi} = \frac{\partial B}{\partial \phi} - \frac{\partial A}{\partial \phi} \theta,$$

where A , θ , and B have the same definition than in (3.9),

$$\begin{aligned}
\frac{\partial A}{\partial \phi} &= \begin{pmatrix} \sum_W \delta & \sum_W x\delta & \sum_W y\delta \\ \sum_W x\delta & \sum_W (x^2 + \frac{\mu g(u_0)}{\alpha})\delta & \sum_W xy\delta \\ \sum_W y\delta & \sum_W xy\delta & \sum_W (y^2 + \frac{\mu g(u_0)}{\alpha})\delta \end{pmatrix} \\
\text{and } \frac{\partial B}{\partial \phi} &= \begin{pmatrix} \sum_W u_0\delta \\ \sum_W (xu_0 + \frac{\mu g(u_0)}{\alpha} \frac{\partial u_0}{\partial x})\delta \\ \sum_W (yu_0 + \frac{\mu g(u_0)}{\alpha} \frac{\partial u_0}{\partial y})\delta \end{pmatrix}.
\end{aligned}$$

It can be noted that this problem expresses the same particularities as (3.9).

As one can see values for the $\frac{\partial a_i(\mathbf{x})}{\partial \phi}$, $\frac{\partial b_i(\mathbf{x})}{\partial \phi}$, and $\frac{\partial c_i(\mathbf{x})}{\partial \phi}$ are proportional to sums over the contour while those for the a_i , b_i , and c_i are proportional to sums over region domains. This suggests that in (I.1) the influence of the *supplementary* term may be small by comparison to that of the *usual* term. A similar assumption has been made by Rosenhahn *et al.* (Rosenhahn *et al.*, 2007).




Figure-A I-1: Energy level in function of the number of iterations during the segmentation of a BrainWeb image with the exact and approximate implementation.

This hypothesis has been verified empirically by comparing the behavior of the exact minimization flow, based on (I.1), and the approximate minimization flow, as given by (3.13). Fig. I-1 shows the progression of the energy while segmenting an image from the BrainWeb dataset with 3% of noise and 40% INU. It appears that the two methods converge smoothly to a very similar energy level. The exact minimization flow converge a bit faster at the beginning, but its computational cost is much higher. In fact, for a four phase implementation, computing the exact flow requires performing roughly twice the number of convolutions and solving three times as many linear systems. The approximate flow is also simpler to implement.

APPENDIX II

SEMI-AUTOMATIC SEGMENTATION OF MAJOR AORTO-PULMONARY COLLATERAL ARTERIES (MAPCAS) FOR IMAGE GUIDED PROCEDURES

David Rivest-Hénault¹, Luc Duong¹, Chantale Lapierre²,
Sylvain Deschênes² and Mohamed Cheriet¹,

¹ Département de génie de la production automatisée, École de Technologie Supérieure,
1100 Notre-Dame Ouest, Montréal, Québec, Canada H3C 1K3

² Sainte-Justine Hospital,
3175 Côte-Ste-Catherine, Montréal, Québec, H3T 1C5, Canada

Paper published in the proceeding of SPIE Medical Imaging 2010

Abstract

Manual segmentation of pre-operative volumetric dataset is generally time consuming and results are subject to large inter-user variabilities. Level-set methods have been proposed to improve segmentation consistency by finding interactively the segmentation boundaries with respect to some priors. However, in thin and elongated structures, such as major aorto-pulmonary collateral arteries (MAPCAs), edge-based level set methods might be subject to flooding whereas region-based level set methods may not be selective enough. The main contribution of this work is to propose a novel expert-guided technique for the segmentation of the aorta and of the attached MAPCAs that is resilient to flooding while keeping the localization properties of an edge-based level set method. In practice, a two stages approach is used. First, the aorta is delineated by using manually inserted seed points at key locations and an automatic segmentation algorithm. The latter includes an intensity likelihood term that prevents leakage of the contour in regions of weak image gradients. Second, the origins of the MAPCAs are identified by using another set of seed points, then the MAPCAs' segmentation boundaries are evolved while being constrained by the aorta segmentation. This prevents the aorta to interfere with the segmentation of the MAPCAs. Our preliminary results are promising and constitute an indication that an accurate segmentation of the aorta and MAPCAs can be obtained with reasonable amount of effort.

Keywords: CTA, segmentation, level-set, MAPCAs, aorta

1 Introduction

In certain severe cases of tetralogy of Fallot or of pulmonary atresia, lung perfusion relies totally or in part on one or more major aorto-pulmonary collateral arteries (MAPCAs), which is not desirable and may require surgical interventions in the months following birth. Examination of those arteries is difficult, however, because they show significant variation in shape, position and number. In addition, pediatric cardiology requires high accuracy: diameter of the aorta may be as small as 7 mm and that of the MAPCAs even smaller. Hopefully, recent computed tomography angiography (CTA) systems now routinely offer a resolution sufficient for the visualization of such structures (Hayabuchi *et al.*, 2008). A 3D geometric model of the structures of interest created from a segmentation of the dataset could expose precious topological information to the experts yielding more confidence in diagnostics and easing the planning of surgical procedures. However, the high precision CTA datasets that are needed tend to be large and tedious to work with. Also, the result of a manual segmentation process may vary greatly in function of the expert skills and of the amount of work involved. On the other hand, due to both the large geometric variability and the small tubular shape of the structure of interest the development of automatic segmentation methods is challenging. In this paper, we presented a workflow that allows the user to control each step of segmentation, and our preliminary work on a method to prevent flooding using a two stages approach (aorta and MAPCAs). This will eventually be incorporated to a dedicated user interface, providing a mean to achieve accurate segmentation of difficult structures.

2 Related work

To our knowledge, this work is among the first that targets vessel malformations for paediatric cardiology, where the anatomic structures of interest are relatively small and are very challenging to segment. Although no work found that focus on MAPCAs segmentation, aorta segmentation received great interest from the community. Recently, Išgum *et al.* (Išgum *et al.*, 2009) proposed an aorta segmentation on CT scans method that is based on the registration of multiple atlases. This method is, however, not applicable to the segmentation of abnormal structures and its application to the segmentation of scans of very young patients may be problematic. Another interesting method for aorta segmentation, but on 4D MR volume, has been published by Zhao *et al.* (Zhao *et al.*, 2009). During its initialization step, this method takes benefit from a fast marching level-set method (Sethian, 1999) and of a vessel enhancement filter (Frangi *et al.*, 1998). Because of the limitation of the vessel enhancement filter at Y junc-



Figure-A II-1: In this example, a weak edge cause an active contour to leak. From left to right: input data and user defined seed points, after 500 iterations, and after 900 iterations.

tion(Qian *et al.*, 2009) and of the importance of the partial volume effect for small structure, this method does not seem to be appropriated for MAPCAs segmentation.

Still, vessel segmentation in general is a very active topic and received a great deal of attention in the last few years. Most relevant to this paper are level-set based methods. By example, van Bemmél *et al.*(van Bemmél *et al.*, 2003) used a level-set method for arteries–veins separation on blood pool agent MRI. Also, Manniesing *et al.*(Manniesing *et al.*, 2006a) proposed a similar level-set method for the segmentation of cerebral vasculature. Both use a level-set function of the form: $\phi_t = -F(1 - \varepsilon\kappa)|\nabla\phi|$, with F representing various external forces combined multiplicatively. The contour evolution will stall if any of the force in F equal 0 or if $(1 - \varepsilon\kappa) = 0$ therefore preventing competition between the forces. As a result, those methods will have difficulties to segment regions that are located far from initialization seeds(Manniesing *et al.*, 2006a) or in presence of spurious details. Also, because those methods use a Gaussian intensity model to represent the background, it is required to mask the bones, instruments and other salient structures as a pre-processing step. Another level-set method for the segmentation of cerebral vasculature, is the CURVES method from Lorigos *et al.*(Lorigo *et al.*, 2001), it will be discussed in the next section. Other forms of level-set based segmentation methods with different behaviour have also been proposed in recent years(Chan and Vese, 2001; Paragios and Deriche, 2002; Kim *et al.*, 2005; Cremers *et al.*, 2007).

3 Method

First, the active contour framework on which our approach is based is introduced in the following paragraphs. After, the methodology that has been used to segment the aorta and attached MAPCAs is described in section 3.2.





Figure-A II-2: Sample evolution of our active contour method (II.4) for the segmentation of the MAPCAs in dataset #2. The aorta was pre-segmented and added to the volume rendering for clarity. From left to right: at initialization and after 25, 50, 100, 150, 200 and 250 iterations.

3.1 Active contour framework

Our segmentation algorithms are inspired from the CURVES method (Lorigo *et al.*, 2001) from which they inherit the ability to regularize 1D curves in 3D. This level set based method assumes that the best positions for the boundaries are indicated by strong image gradients, and in this respect is similar to the geodesic active contour method (Caselles *et al.*, 1997). Within the level set framework, a set of boundaries is represented implicitly by the zero level of a 3D function $\phi(\mathbf{x}) : \Omega \subset \mathbb{R}^3 \mapsto \mathbb{R}$. Regions where $\phi(\mathbf{x}) < 0$ indicate voxels belonging to the structure of interest. Starting from an arbitrary surface, CURVES achieves a segmentation by evolving ϕ by respect to an artificial time variable, t :

$$\phi_t = \frac{\partial \phi(\mathbf{x}, t)}{t} = \delta_\epsilon(\phi) \left(\lambda(\nabla \phi(\mathbf{x}, t), \nabla^2 \phi(\mathbf{x}, t)) + \rho \frac{g'}{g} \nabla \phi(\mathbf{x}, t) \cdot \mathbf{H} \frac{\nabla \mathbf{I}(\mathbf{x})}{|\nabla \mathbf{I}(\mathbf{x})|} + \gamma g(|\nabla \mathbf{I}(\mathbf{x})|) \right), \quad (\text{II.1})$$

where $\mathbf{I}(\mathbf{x})$ represents the intensity of the voxel located in \mathbf{x} , \mathbf{H} is the Hessian matrix and $\lambda(\nabla \phi(\mathbf{x}, t), \nabla^2 \phi(\mathbf{x}, t))$ is defined as the smallest principal curvature of $\phi(\mathbf{x})$. Also, $g(|\nabla \mathbf{I}|)$ is a monotonically decreasing edge indicator function such that $g(|\nabla \mathbf{I}|) \rightarrow 0$ as $|\nabla \mathbf{I}| \rightarrow \infty$. A typical choice is $g(z) = \frac{1}{1+z^2}$ (Caselles *et al.*, 1997). The regularized Delta function $\delta_\epsilon(\phi)$ limits the evolution around the zero level of ϕ . The first term of (II.1) is a regularization term that tends to smooth the boundaries, the second is a data advection term that attracts the boundaries toward the edges and the third term is a balloon force that tends to inflate the volume¹.

¹It is worth noting that although the balloon term is not part of the CURVES formulation (Lorigo *et al.*, 2001), this term does appear in the original geodesic active contour formulation (Caselles *et al.*, 1997) on which CURVES is based.

Since weak edges may cause this algorithm to fail to delineate correctly the structures, as shown on Figure II-1, an estimation of the expected voxel intensity is considered. This is taken into account by replacing the balloon term by a new one in the boundaries evolution equation, as follows:

$$\phi_t = \delta_\varepsilon(\phi) \left(\lambda(\nabla\phi(\mathbf{x},t), \nabla^2\phi(\mathbf{x},t)) + \rho \frac{g'}{g} \nabla\phi(\mathbf{x},t) \cdot \mathbf{H} \frac{\nabla\mathbf{I}(\mathbf{x})}{|\nabla\mathbf{I}(\mathbf{x})|} - \gamma g(|\nabla\mathbf{I}(\mathbf{x})|) f(\mathbf{I}(\mathbf{x})) \right). \quad (\text{II.2})$$

The purpose of this new term is to attract the moving boundary toward regions that have voxel intensities similar to the intensities of the voxels located inside the boundaries and to repulse it otherwise. In this work, f is defined as:

$$f(z) = \exp\left(-\frac{(z - \bar{I})^2}{s^2}\right).$$

Here, s is a parameter that accounts for the standard deviation of the intensity and \bar{I} is a local estimation of the intensity of the structure to segment and is computed from the voxels lying inside the boundaries at initialization, i.e. where $\phi_{t=0}(\mathbf{x}) < 0$, using local weighted average operations (Rivest-Hénault and Cheriet, 2007; Li *et al.*, 2008b; Brox and Cremers, 2009):

$$\bar{I}(\mathbf{x}) = \frac{K_\sigma * (H_\varepsilon(\phi)I)}{K_\sigma * H_\varepsilon(\phi)}, \text{ with } H_\varepsilon(z) = \begin{cases} 0 & z < -\varepsilon \\ \frac{1}{2} + \frac{1}{2} \sin(\pi z / (2\varepsilon)) & |z| \leq \varepsilon \\ 1 & z > \varepsilon \end{cases}, \quad (\text{II.3})$$

K_σ is a Gaussian kernel with standard deviation σ and $*$ denotes the convolution operator. The parameter ε of the $H_\varepsilon(z)$ function account for the ambiguity of the contour position and is set to the distance separating two adjacent voxels. Equation (II.2) requires that \bar{I} be defined anywhere on the image domain. Therefore, the estimation for \bar{I} must be extended to regions where $\phi \geq 0$. A simple method is to use the values of the closest point where $\phi < 0$. Finally, ρ and γ are free parameters.

It can also be useful to include an additional advection force that would push the boundaries away from a certain object. This can be done by adding a new force to (II.2), as follows:

$$\phi_t = \delta_\varepsilon(\phi) \left(\lambda(\nabla\phi(\mathbf{x},t), \nabla^2\phi(\mathbf{x},t)) + \rho \frac{g'}{g} \nabla\phi(\mathbf{x},t) \cdot \mathbf{H} \frac{\nabla\mathbf{I}}{|\nabla\mathbf{I}|} - \gamma g(\mathbf{I}) f(\mathbf{I}) - \beta h(d(\mathbf{x})) \frac{\nabla\phi(\mathbf{x},t)}{|\nabla\phi(\mathbf{x},t)|} \cdot \nabla d(\mathbf{x}) \right). \quad (\text{II.4})$$

Here, $d(\mathbf{x}) : \Omega \mapsto \mathbb{R}$ is the distance transform of a certain object in Ω , $h(d)$ is a function that limits the reach of this advection force and β is a free parameter. In this work, this force is used to push the MAPCAs boundaries away from the aorta. Also, we set $h(d) = \exp(-d^2/k)$. For an example evolution of an active contour by using (II.4), see Figure II-2.

3.2 Description of our approach

It is desired to segment the aorta and the attached MAPCAs in CTA volumes by finding the surface that best separates those structures from the volumetric dataset. The proposed method is composed of two semi-automatic stages. In the first stage, a) several key points are placed inside the aorta by a user by clicking on MPR views of the volume, and b) the aorta is delineated automatically by using a segmentation algorithm described below. In the second stage, a) seed points indicating the origin of the MAPCAs are placed close to the segmented aorta by the user, and b) the final segmentation of the aorta and of the attached MAPCAs is obtained automatically by using the second segmentation algorithm described below. Also, it is worth noting that, as a pre-processing step, spherical masks were manually placed at places where the aorta or a MAPCA entered the heart in order to prevent the active contour to fill the heart chambers.

In the first stage, a user is required to place small spheres of radius r inside the aorta. After, the fast marching method (Sethian, 1999) (FMM) is used to create an initial level set function ϕ . Voxels inside the spheres are given negative values and those outside positive values. The level set function ϕ is then evolved using (II.2) until the maximum boundary displacement is smaller than a certain value or until stopped by the user. This process results in a segmentation of the aorta, as shown on Figure II-3.

In the second stage, the user place small spheres, also of radius r , indicating the junction of the MAPCAs and of the aorta. Then, a level set function is created as described earlier and the segmentation of the MAPCAs is produced by using (II.4). For an example, see Figure II-4. This algorithm takes benefit from the volume describing the aorta in two ways: 1) starting from the aorta, the FFM is used to create an advection field that is used to push the MAPCAs segmentation boundaries away from the aorta; and 2) the aorta volume is subtracted from the domain of the MAPCAs level set function. This prevents the aorta voxel to interfere with the MAPCAs segmentation.




Figure-A II-3: Aorta segmentation in dataset #1. From left to right: 1) an MPR view of the input data with seed points and the resulting aorta segmentation highlighted, 2–3) a volume rendering view of the segmented aorta viewed from 2) the anteroposterior direction, and 3) from the mediolateral direction.



Figure-A II-4: MAPCAs segmentation in dataset #1. From left to right: 1) an MPR view of the input data with seed points and the resulting aorta segmentation highlighted, 2–3) a volume rendering view of the segmented aorta and attached MAPCAs viewed from 2) the anteroposterior direction, and 3) the mediolateral direction.

4 Experimentation and results

The algorithms were implemented in C++ using ITK² and the XipBuilder framework³. For a $256 \times 256 \times 110$ voxels volume, computation times were evaluated at around 30 sec for the aorta segmentation and around 1 min for the MAPCAs segmentation using our single threaded implementation on a AMD X2 6000+ CPU.

The parameters of our method must be selected based on prior knowledge or hand picked by trial and errors. The radius of the seeds points, r was set to 1.5mm, the approximate radius

²Available from the Insight Software Consortium at <http://www.itk.org>

³Available from Siemens at <https://collab01a.scr.siemens.com/xipwiki/index.php/XipBuilder>




Figure-A II-5: Comparison between the segmentation results obtained by the proposed method, images (a)-(c) and (g)-(i), and those obtained by using the CURVE equation with the propagation term (II.1), as described in section 4, images (d)-(f) and (j)-(l). Top row: dataset #1, bottom row: dataset #2. The gap indicated by an arrow in (h) and (k) is caused by the presence of a catheter.




Figure-A II-6: Comparison between the segmented pre-operative 3D data and a frame from a 2D+t angiographic acquisition for dataset #1. Those images are from the same patient than in Figures II-3 and II-4.

of the smallest MAPCA. The parameters ρ , γ and s of (II.2) and (II.4) depend on the contrast of the dataset. In our case, all voxel intensities were normalized between 0 and 1 and the parameters were set as follows: $\rho = 5$, $\gamma = 2$ and $s \in [0.04, 0.05]$. Finally, in (II.4) the parameter β was set to $0.1 \cdot \rho$ and k was set to 10mm, which corresponds approximately to the diameter of the aorta.

An example segmentation of the aorta can be seen in Figure II-3 and an example segmentation of the four MAPCAs attached to this aorta is shown in Figure II-4. For comparison purpose, a 2D angiographic view of the same structures is presented on Figure II-6.

In order to illustrate the benefits of our approach, the same two datasets have been segmented by using the CURVES level-set equation (II.1) in place of (II.2) and (II.4) in the methodology described in section 3.2. The same seed points were also used to initialize both steps. It should be noted that this differs from the methodology introduced by Lorigo *et al.* (Lorigo *et al.*, 2001). The parameters of (II.1) have been selected in order to maximize the extend of parts of the MAPCAs that could be segmented while keeping leakage at a minimum. Those comparisons are presented on Figure II-5. Four MAPCAs originating from the descending aorta were identified in dataset #1 (top row, in Figure II-5). By using the proposed method, all four appeared to be correctly segmented. Also, the contour was able to propagate itself far into the lungs' vessel. Most of the segmented vessels were well separated from the others. When using the modified CURVES equation, it was also possible to segment the four MAPCAs in region close to the aorta, but the length of the vessels that was segmented was much smaller. In this case, the level-set propagation was stop by hand, but more iterations caused to contour to leak. The dataset #2 (bottom row, in Figure II-5) was more difficult to segment because the scan was less contrasted than for dataset #1. In this case, two MAPCAs also originating from the descending aorta were identified. By using the proposed method, it was possible to segment the two MAPCAs. On the contrary, it was not possible to obtain satisfactory results by using the modified CURVES method. As it can be seen on Figure II-5-(j), the contour started to leak close to the aorta.

4.1 Discussion

The segmentation of MAPCAs on CTA is challenging. Those vessels are of very small size, diameters of 1-3mm being frequent. Because of this, their apparent intensity in CTA is affected by partial volume effect. Moreover, the origin, size and shape of the MAPCAs vary greatly from one case to the other. Our preliminary results indicate that our technique is able to deliver accurate segmentation results for this difficult task. In all cases, the active contours were less

subject to leakage when using (II.2) or (II.4) instead of (II.1). That can be explained by our driving term, $f(I(\mathbf{x}))$, that limits the propagation of the contour to regions that have intensities similar to that of the seeds points. As a result, it is possible to perform more iterations without experimenting leakage. In fact, when using (II.2) or (II.4) the propagation of the contour stopped by itself. Also, it is worth noting that our model is able to take benefit from regional intensity information for the structures of interest while it does not required a modelization of the background. Therefore, it is not needed to mask the bone or the instruments prior to the segmentation. As for our two-step approach, one of its advantages is that the segmentation of the aorta in the first step permits the creation of an advection field that can be used to support the segmentation of the MAPCAs. This has the potential to help to segment very small vessels. The evaluation of the benefit of this aspect is to be done in a future communication. Another advantage is that our approach allows to use different sets of parameters for the segmentation of the aorta and for the segmentation of the MAPCAs. However, we did not take advantage of this possibility in our experiments.

In future work, this segmentation method will be integrated in an image-guided navigation framework. We plan to take advantage of the segmented aorta and MAPCAs to create patient-specific 3D models that will be registered with 2D angiographic sequences, as Figure II-6 suggests. We hope that such registration would provide valuable information to complement the 2D angiography during surgical procedures by offering depth indications and temporal consistency.

5 Conclusion

The level set method is an interesting framework for the integration of various segmentation principles that can lead to high quality segmentation results. However, in the case of the segmentation of thin elongated structures in CTA volumes of paediatric patients, classical segmentation methods may not be directly usable due to practical problems. A new two stages semi-automatic technique for the segmentation of the aorta and attached MAPCAs has been presented. In the first stage, the aorta is segmented by using seed points placed by a user and a new level set formulation that includes local estimations of the intensity of the structures to segment. In the second stage, the boundary delineating the aorta is frozen and the attached MAPCAs are segmented. A particularity of this algorithm is that it takes advantage of an advection force that tends to push the MAPCAs boundaries away from the aorta. Preliminary results indicate that our technique is able to deliver accurate and repeatable segmentation results for the aorta and attached MAPCAs.

BIBLIOGRAPHY

- Iman Aganj, Christophe Lenglet, and Guillermo Sapiro. 2010. ODF maxima extraction in spherical harmonic representation via analytical search space reduction. Jiang, T., Nasir Navab, Josien Pluim, and Max Viergever, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2010*, volume 6362 of *Lecture Notes in Computer Science*, p. 84–91. Springer Berlin / Heidelberg. ISBN 978-3-642-15744-8.
- AHA News. 01 2008. “Heart and stroke death rates steadily decline; risks still too high”. [Online; accessed 2-April-2008]. <<http://www.americanheart.org/presenter.jhtml?identifier=3053235>>.
- Ayelet Akselrod-Ballin, Meirav Galun, Moshe John Gomori, Ronen Basri, and Achi Brandt. 2006. “Atlas Guided Identification of Brain Structures by Combining 3D Segmentation and SVM Classification.”. In *MICCAI (2)*.
- Elsa D. Angelini, Ting Song, Brett D. Mensh, and Andrew F. Laine. 2007. “Brain MRI Segmentation with Multiphase Minimal Partitioning: A Comparative Study”. *Int J Biomed Imaging*, vol. 10526, p. 1–15.
- Souha Aouadi and Laurent Sarry. 2008. “Accurate and precise 2D-3D registration based on X-ray intensity”. *Comput. Vis. Image Underst.*, vol. 110, n. 1, p. 134–151.
- Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. 2006. “Log-Euclidean Metrics for Fast and Simple Calculus on Diffusion Tensors”. *Magnetic Resonance in Medicine*, vol. 56, p. 411–421.
- Peter J. Basser and Carlo Pierpaoli. 1996. “Microstructural and Physiological Features of Tissues Elucidated by Quantitative-Diffusion-Tensor MRI”. *Journal of Magnetic Resonance, Series B*, vol. 111, n. 3, p. 209 – 219.
- Pierre-Louis Bazin and Dzung L. Pham. 2007. “Topology-Preserving Tissue Classification of Magnetic Resonance Brain Images”. *IEEE Trans. Med. Imaging*, vol. 26, n. 1, p. 487–496.
- Boubakeur Belaroussi, Julien Milles, Sabin Carme, Yue Min Zhu, and Hugues Benoit-Cattin. 2006. “Intensity non-uniformity correction in MRI: Existing methods and their validation”. *Medical Image Analysis*, vol. 10, n. 2, p. 234–246.
- Ismail Ben Ayed, Amar Mitiche, and Ziad Belhadj. 2006. “Polarimetric Image Segmentation via Maximum-Likelihood Approximation and Efficient Multiphase Level-Sets”. *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, p. 1493–1500.
- Said Benameur, Max Mignotte, Stefan Parent, Hubert Labelle, Wafa Skalli, and Jacques de Guise. 2003. “3D/2D registration and segmentation of scoliotic vertebrae using statistical models”. *Computerized Medical Imaging and Graphics*, vol. 27, p. 321–337.

- F. Benmansour, L. D. Cohen, Wai Kong Law, and Albert Chi-Shing Chung. 2009. "Tubular anisotropy for 2D vessel segmentation". In *CVPR 2009*.
- S. Berlemont and J.-C. Olivo-Marin. jan. 2010. "Combining Local Filtering and Multiscale Analysis for Edge, Ridge, and Curvilinear Objects Detection". *Image Processing, IEEE Transactions on*, vol. 19, n. 1, p. 74–84.
- Wolfgang Birkfellner, R. Seemann, Michael Figl, Johann Hummel, C. Ede, P. Homolka, X. Yang, P. Niederer, and Helmar Bergmann. 2005. "Wobbled splatting — a fast perspective volume rendering method for simulation of X-ray images from CT". *Physics in medicine and biology*, vol. 50, p. 73–84.
- Wolfgang Birkfellner, Michael Figl, Christelle Gendrin, Johann Hummel, Shuo Dong, and Helmar Bergmann. 2009. "Stochastic rank correlation: A robust merit function for 2D/3D registration of image data obtained at different energies". *Medical Physic*, vol. 36, p. 3420–3428.
- C. Blondel, G. Malandain, R. Vaillant, and N. Ayache. may 2006. "Reconstruction of coronary arteries from a single rotational X-ray projection sequence". *Medical Imaging, IEEE Transactions on*, vol. 25, n. 5, p. 653–663.
- Jonathan Boisvert, Farida Cheriet, Xavier Pennec, Hubert Labelle, and Nicholas Ayache. 2008. "Geometric Variability of the Scoliotic Spine Using Statistics on Articulated Shape Models". *IEEE TMI*, vol. 27, n. 4, p. 557–568.
- F.L. Bookstein. 1989. "Principal Warps: Thin-Plate Splines and the Decomposition of Deformations". *IEEE PAMI*, vol. 11, n. 6, p. 567–585.
- S. Bouattour, R. Arndt, and D. Paulus. 2005. "4D Reconstruction of Coronary Arteries from Monoplane Angiograms". In *CAIP*.
- Sylvain Bouix, Kaleem Siddiqi, and Allen Tannenbaum. 2005. "Flux driven automatic centerline extraction". *Medical Image Analysis*, vol. 9, n. 3, p. 209–221.
- Yuri Boykov, Olga Veksler, and Ramin Zabih. 2001. "Fast Approximate Energy Minimization via Graph Cuts". *IEEE PAMI*, vol. 23, p. 1222–1239.
- Marco Bozzali, Mara Cercignani, and Carlo Caltagirone. 2008. "Brain volumetrics to investigate aging and the principal forms of degenerative cognitive decline: a brief review". *Magnetic Resonance Imaging*, vol. 26, p. 1065–1070.
- R.P. Brent, 1973. *Algorithms for Minimization without Derivatives*. Englewood Cliffs, NJ : Prentice-Hall.
- T. Brox and D. Cremers. 2009. "On Local Region Models and a Statistical Interpretation of the Piecewise Smooth Mumford-Shah Functional". *Int J Comput Vis*, vol. 84, p. 184–193.
- Thomas Brox. 2005. "From Pixels to Regions: Partial Differential Equations in Image Analysis". PhD thesis, Saarbrücken.

- Thomas Brox and Joachim Weickert. 2004. "Level Set Based Image Segmentation with Multiple Regions". In *Pattern Recognition, 26th DAGM Symposium, Tübingen, Germany, August 30 - September 1, 2004. Proceedings*. p. 415–423. Springer-Verlag.
- Thomas Brox, Bodo Rosenhahn, Juergen Gall, and Daniel Cremers. 2010. "Combined Region- and Motion-based 3D Tracking of Rigid and Articulated Objects". *IEEE PAMI*, vol. 32, p. 402–415.
- T.D. Bui, S. Gao, and Q.H. Zhang. 2005. "A Generalized Mumford–Shah Model for Roof-Edge Detection". *ICIP*, vol. 2, p. II-1214-1217.
- Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. 1997. "Geodesic Active Contours". *Int J Comput Vis*, vol. 22, p. 61–79.
- Tony F. Chan and Luminita Vese. February 2001. "Active Contours Without Edge". *IEEE Trans. Image Process.*, vol. Volume 10, n. 2, p. 266-277.
- L. D. Cohen and R. Kimmel. 1997. "Global minimum for active contour models: a minimal path approach". *IJCV*, vol. 24, p. 57–78.
- DL Collins, AP Zijdenbos, V Kollokian, JG Sled, NJ Kabani, CJ Holmes, and AC Evans. 1998. "Design and Construction of a Realistic Digital Brain Phantom". *IEEE Tran. Med. Imag.*, vol. 17, p. 463–468.
- D. Comaniciu and P. Meer. 2002. "Mean shift: a robust approach toward feature space analysis". *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, n. 5, p. 603–619.
- E Cosman. 2000. "Rigid registration of MR and biplanar fluoroscopy". Master's thesis, Massachusetts Institute of technology.
- C. Couprie, L. Grady, L. Najman, and H. Talbot. july 2011. "Power Watershed: A Unifying Graph-Based Optimization Framework". *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, n. 7, p. 1384–1399.
- Daniel Cremers, Oliver Fluck, Mikael Rousson, and Shmuel Aharon. 2007. "A Probabilistic Level Set Formulation for Interactive Organ Segmentation". In *SPIE Medical Imaging*.
- M.B. Cuadra, L. Cammoun, T. Butz, O. Cuisenaire, and J.-P. Thiran. 2005. "Comparison and validation of tissue modelization and statistical classification methods in T1-weighted MR brain images". *IEEE Tran. Med. Imag.*, vol. 24, n. 12, p. 1548-1565.
- Sean C.L. Deoni, Brian K. Rutt, and Terry M. Peters. 2006. "Synthetic T1-weighted brain image generation with incorporated coil intensity correction using DESPOT1". *Magnetic Resonance Imaging*, vol. 24, n. 9, p. 1241–1248.
- M. Descoteaux, R. Deriche, T.R. Knosche, and A. Anwender. feb. 2009. "Deterministic and Probabilistic Tractography Based on Complex Fibre Orientation Distributions". *Medical Imaging, IEEE Transactions on*, vol. 28, n. 2, p. 269 -286.

- Maxime Descoteaux and Rachid Deriche. 2008. "High Angular Resolution Diffusion MRI Segmentation Using Region-Based Statistical Surface Evolution". *Journal of Mathematical Imaging and Vision*, vol. 33, n. 2, p. 239–252.
- Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, and Rachid Deriche. 2007. "Regularized, fast, and robust analytical Q-ball imaging.". *Magnetic Resonance in Medicine*, vol. 58, n. 3, p. 497–510.
- Richard O. Duda, Peter E. Hart, and David G. Stork, 2000. *Pattern Classification*. éd. 2nd Edition. Wiley-Interscience.
- L. Duong, R. Liao, H. Sundar, B. Tailhades, and C. Xu. 2009. "Curve-based 2D-3D registration of coronary vessels for image guided procedure". In *Proceedings of SPIE Medical Imaging*.
- R. Fabbri, L. da F. Costa, J. C. Torelli, and O. M. Bruno. 2008. "2D Euclidean Distance Transform Algorithms: A Comparative Survey". *ACM Computing Surveys*, vol. 40, n. 1, p. 1–44.
- Jacques Feldmar, Nicholas Ayache, and Fabienne Betting. 1994. *3d-2d projective registration of free-form curves and surfaces*. Technical report. INRIA.
- Jacques Feldmar, Nicholas Ayache, and Fabienne Betting. 1995. "3D-2D projective registration of free-form curves and surfaces". In *Fifth International Conference on Computer Vision (ICCV'95)*.
- Jacques Feldmar, Nicholas Ayache, and Fabienne Betting. 1997. "3D-2D projective registration of free-form curves and surfaces". *Computer Vision and Image Understanding*, vol. 65, p. 403-424.
- Ferreira. 2007. "A Dirichlet process mixture model for brain MRI tissue classification". *Medical Image Analysis*, vol. 11, p. 169-182.
- David J. Field. 1987. "Relations between the statistics of natural images and the response properties of cortical cells". *J. Opt. Soc. Am. A*, vol. 4, p. 2379–2394.
- Charles Florin, Nikos Paragios, and Jim Williams. 2005. "Particle filters, a quasi-monte carlo solution for segmentation of coronaries". In *MICCAI'05 Proceedings of the 8th international conference on Medical Image Computing and Computer-Assisted Intervention*.
- Fondation des maladies du coeur. 2008. "Statistiques". <http://www.who.int/cardiovascular_diseases/en/>. [Online, accessed 17-Nov-2008].
- Alejandro F. Frangi, Wiro J. Niessen, Koen L. Vincken, and Max A. Viergever. 1998. "Multiscale vessel enhancement filtering". In *MICCAI'98*.
- Alejandro F. Frangi, Wiro J. Niessen, Paul J. Nederkoorn, Jeannette Bakker, Willem P.Th.M. Mali, and Max A. Viergever. 2001. "Quantitative Analysis of Vascular Morphology From 3D MR Angiograms: In Vitro and In Vivo Results". *Magnetic Resonance in Medicine*, vol. 45, p. 311–322.

- Lawrence R. Frank. 2002. "Characterization of anisotropy in high angular resolution diffusion-weighted MRI". *Magn. Reson. Med*, vol. 47, p. 1083–1099.
- Ola Friman, Milo Hindennach, Caroline Kühnel, and Heinz-Otto Peitgen. 2010. "Multiple hypothesis template tracking of small 3D vessel structures". *Medical Image Analysis*, vol. 14, n. 2, p. 160 - 171.
- Christopher L.F. Gade and S. Chiu Wong, 2006. *SIS 2006 Yearbook*, chapter CHRONIC TOTAL OCCLUSIONS, p. 51–58.
- Song Gao and Tien D. Bui. 2004. "A new images segmentation and smoothing model". In *ISBI*.
- Song Gao and Tien D. Bui. 2005. "Image Segmentation and Selective Smoothing by Using Mumford–Shah Model". *IEEE Trans. Image Process.*, vol. 14, n. 10, p. 1537-1549.
- Y. Gao, A. Tannenbaum, and R. Kikinis. 2010. "Simultaneous Multi-Object Segmentation using Local Robust Statistics and Contour Interaction". In *Workshop on Medical Computer Vision, MICCAI 2010*.
- E. J. Garboczi. 2002. "Three-Dimensional Mathematical Analysis of Particle Shape Using X-Ray Tomography And Spherical Harmonics: Application To Aggregates Used In Concrete". *Cement and Concrete Research*, vol. 32, n. 10, p. 1621–1638.
- Stuart Geman and Donald Geman. 1984. "Stochastic Relaxation, Gibbs Distribution and the Bayesian Restoration of Images". *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, p. 721-741.
- José Gomes and Olivier Faugeras. 2000. "Reconciling Distance Functions and Level Sets". *Journal of Visual Communication and Image Representation*, vol. 11, p. 209–223.
- Leo Grady. 2006. "Random Walks for Image Segmentation". *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, p. 1-17.
- J. Aaron Grantham, Steven P. Marso, John Spertus, John House, David R. Holmes JR, and Barry D. Rutherford. 2009. "Chronic Total Occlusion Angioplasty in the United States". *JACC: CARDIOVASCULAR INTERVENTIONS*, vol. 2, n. 6, p. 479–486.
- M Grass, R Koppe, E Klotz, R Proksa, M.H Kuhn, H Aerts, J Op de Beek, and R Kemkers. 1999. "Three-dimensional reconstruction of high contrast objects using C-arm image intensifier projection data". *Computerized Medical Imaging and Graphics*, vol. 23, n. 6, p. 311 - 321.
- Henry Gray, 1918. *Anatomy of the Human Body*. Lea & Febiger.
- H. Greenspan, A. Ruf, and J. Goldberger. 2006. "Constrained Gaussian mixture model framework for automatic segmentation of MR brain images". *IEEE Trans. Med. Imag.*, vol. 25, p. 1233–1245.

- D. Greig, B. Porteous, and A.. Seheult. 1989. "Exact maximum a posteriori estimation for binary images". *Journal of the Royal Statistical Society, Serie B*, vol. 51, p. 271–279.
- M. Groher, D. Zikic, and N. Navab. june 2009. "Deformable 2D-3D Registration of Vascular Structures in a One View Scenario". *IEEE TMI*, vol. 28, n. 6, p. 847–860.
- Hákon Gudbjartsson and Samuel Patz. 1995. "The Rician Distribution of Noisy MRI Data". *Magn Reson Med.*, vol. 34, p. 910–914.
- M. Gülsün and Hüseyin Tek. 2008. "Robust Vessel Tree Modeling". In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2008*. p. 602-611.
- L.O. Hall, A.M. Bensaid, L.P. Clarke, R.P. Velthuizen, M.S. Silbiger, and J.C. Bezdek. sep 1992. "A comparison of neural network and fuzzy clustering techniques in segmenting magnetic resonance images of the brain". *Neural Networks, IEEE Transactions on*, vol. 3, n. 5, p. 672–682.
- R. I. Hartley and A. Zisserman, 2004. *Multiple View Geometry in Computer Vision*. éd. 2nd. Cambridge University Press, ISBN: 0521540518.
- Yasunobu Hayabuchi, Miki Inoue, Noriko Watanabe, Miho Sakata, Manal Mohamed Helmy Nabo, Tetsuya Kitagawa, Takashi Kitaichi, and Shoji Kagami. 2008. "Assessment of systemic-pulmonary collateral arteries in children with cyanotic congenital heart disease using multidetector-row computed tomography: Comparison with conventional angiography". *International Journal of Cardiology*, vol. IN PRESS, p. 1–6.
- John H. Hipwell, Graeme P. Penney, Robert A. McLaughlin, Kawal Rhode, Paul Summers, Tim C. Cox, James V. Byrne, J. Alison Noble, and David J. Hawkes. 2003. "Intensity-Based 2-D–3-D Registration of Cerebral Angiograms". *IEEE TMI*, vol. 22, p. 1417–1426.
- Zhang Hongmei and Wan Mingxi. 2006. "Improved Mumford–Shah Functional for Coupled Edge-Preserving Regularization and Image Segmentation". *EURASIP Journal on Applied Signal Processing*, vol. 1, p. 1-9.
- Aviv Hurvitz and Leo Joskowicz. 2008. "Registration of a CT-like atlas to fluoroscopic X-ray images using intensity correspondences". *International Journal of Computer Assisted Radiology and Surgery*, vol. 3, p. 493–504.
- Ivana Išgum, Marius Staring, Annemarieke Rutten, Mathias Prokop, Max A. Viergever, and Bram van Ginneken. 2009. "Multi-Atlas-Based Segmentation With Local Decision Fusion—Application to Cardiac and Aortic Segmentation in CT Scans". *IEEE Trans. Med. Imag.*, vol. 28, n. 7, p. 1000–1010.
- J. R. Jimenez-Alaniz, V. Medina-Banuelos, and Oscar Yáñez-Suárez. 2006. "Data-driven brain MRI segmentation supported on edge confidence and a priori tissue information". *IEEE Trans. Med. Imaging*, vol. 25, n. 1, p. 74-83.

- Steven G. Johnson. 2012. “The NLOpt nonlinear-optimization package”. <<http://ab-initio.mit.edu/nlopt>>.
- B. Johnston, M. S. Atkins, B. Mackiewicz, and M. Anderson. 1996. “Segmentation of Multiple Sclerosis Lesions in Intensity Corrected Multispectral MRI”. *IEEE TMI*, vol. 15, n. 2, p. 154–169.
- D. R. Jones, C. D. Perttunen, and B. E. Stuckmann. 1993. “Lipschitzian optimization without the lipschitz constant”. *J. Optimization Theory and Applications*, vol. 79, p. 157.
- David Rivest-Hénault Julien Couet and, Joaquim Miró, Chantale Lapierre, Luc Duong, and Mohamed Cheriet. 2012. “Intensity-based 3D/2D registration for percutaneous intervention of major aorto-pulmonary collateral arteries”. In *Spie Medical Imaging*.
- M. Kass, A. Witkin, and D. Terzopoulos. 1988. “Snakes: Active contour models.”. *Int. J. of Computer Vision*, vol. 1, p. 321–331.
- DN Kennedy, PA Filipek, and VS Caviness. 1989. “Anatomic segmentation and volumetric calculations in nuclear magnetic resonance imaging”. *IEEE Tran. Med. Imag.*, vol. 8, p. 1-7.
- E. Kerrien, M o. Berger, E. Maurincomme, L. Launay, and R. Vaillant. 1999. “Fully automatic 3D/2D subtracted angiography registration”. In *Medical Image Computing and ComputerAssisted Intervention — MICCAI’99*. p. 664–671. Springer-Verlag.
- Ali Khamene, Peter Bloch, Wolfgang Wein, Michelle Svatos, and Frank Sauer. 2006. “Automatic registration of portal images and volumetric CT for patient positioning in radiation therapy”. *Medical Image Analysis*, vol. 10, p. 96–112.
- Junmo Kim, John W. Fisher III, Anthony Yezzi, Müjdat Çetin, and Alan S. Willsky. August 2002. “Nonparametric Methods for Image Segmentation Using Information Theory and Curve Evolution”. In *IEEE International Conference on Image Processing*. p. 1169-1186.
- Junmo Kim, John W. Fisher III, Anthony Yezzi, Müjdat Çetin, and Alan S. Willsky. October 2005. “A Nonparametric Statistical Method for Image Segmentation Using Information Theory and Curve Evolution”. *IEEE Trans. Image Process.*, vol. 14, n. 10, p. 1486-1502.
- Ron Kimmel, 2003. *Fast Edge Integration*, chapter 4, p. 59-77. Springer.
- Vladimir Kolmogorov. 2004. “What Energy Functions Can Be Minimized via Graph Cuts?”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, p. 147-159.
- Karl Krissian, Xunlei Wu, and Vincent Luboz, 2006. *Medicine Meets Virtual Reality 14: Accelerating Change in Healthcare: Next Medical Toolkit*, chapter Smooth Vasculture Reconstruction with Circular and Elliptic Cross Sections, p. 273-278. IOS Press.
- Hsiang-Ching Kung, Donna L. Hoyert, Jiaquan Xu, and Sherry L. Murphy. 2008. “Deaths: Final Data for 2005”. *National Vital Statistics Reports*, vol. 56, p. 1–66.

- K. K. Lau and Albert C. S. Chung. 2006. "A global optimization strategy for 3D–2D registration of vascular images". In *The 17th British Machine Vision Conference, (BMVC 06)*.
- M. Law and A. Chung. march 2009. "Efficient Implementation for Spherical Flux Computation and Its Application to Vascular Segmentation". *Image Processing, IEEE Transactions on*, vol. 18, n. 3, p. 596 -612.
- M. W. K. Law and A. C. S. Chung. 2008. "Three dimensional curvilinear structure detection using optimally oriented flux". In *ECCV*. p. 268–382.
- Max W. K. Law and Albert C. S. Chung. 2010. "An oriented flux symmetry based active contour model for three dimensional vessel segmentation". In *Proceedings of the 11th European conference on computer vision conference on Computer vision: Part III*. (Berlin, Heidelberg 2010), p. 720–734. Springer-Verlag.
- Suk-Ho Lee and Jin Keun Seo. 2006. "Level Set-Based Bimodal Segmentation With Stationary Global Minimum". *IEEE transactions on image processing*, vol. 15, p. 2843-2852.
- C. Lemaitre, M. Perdoch, A. Rahmoune, J. Matas, and J. Miteran. 2011. "Detection and matching of curvilinear structures". *Pattern Recognition*, vol. 44, n. 7, p. 1514 - 1527.
- David Lesage, Elsa D. Angelini, Isabelle Bloch, and Gareth Funka-Lea. 2009a. "A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes". *Medical Image Analysis*, vol. 13, n. 6, p. 819–845.
- David Lesage, Elsa D. Angelini, Gareth Funka-Lea, and Isabelle Bloch. 2009b. "Bayesian maximal paths for coronary artery segmentation from 3D CT angiograms". In *MICCAI*.
- Hava Lester and Simon R. Arridge. 1999. "A survey of hierarchical non-linear medical image registration". *Pattern Recognition*, vol. 32, n. 1, p. 129–149.
- C. Li, R. Huang, Z. Ding, C. Gatenby, D. N. Metaxas, and J. C. Gore. 2008a. "A Variational Level Set Approach to Segmentation and Bias Correction of Images with Intensity Inhomogeneity". In *MICCAI*.
- Chunming Li, Chenyang Xu, Changfeng Gui, and Martin D. Fox. 2005a. "Level Set Evolution without Re-Initialization: A New Variational Formulation". In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1*. (Washington, DC, USA 2005), p. 430–436. IEEE Computer Society.
- Chunming Li, Chenyang Xu, Kishori M. Konwar, and Martin D. Fox. 2006. "Fast Distance Preserving Level Set Evolution for Medical Image Segmentation". In *ICARCV*.
- Chunming Li, Chiu-Yen Kao, John C. Gore, and Zhaohua Ding. 2008b. "Minimization of Region-Scalable Fitting Energy for Image Segmentation". *IEEE Trans. Image Processing*, vol. 17, n. 10, p. 1940–1949.

- Hua Li, Anthony Yezzi, and Laurent D. Cohen. 2005b. “Fast 3D Brain Segmentation Using Dual-Front Active Contours with Optional User-Interaction”. In *Computer Vision for Biomedical Image Applications*. p. 335–345. Springer Berlin / Heidelberg.
- Rui Liao, Christoph Guetter, Chenyang Xu, Yiyong Sun, Ali Khamene, and Frank Sauer. 2006. “Learning-Based 2d/3d rigid registration using jensen-shannon divergence for image-guided surgery”. In *Proceedings of the Third international conference on Medical Imaging and Augmented Reality*. (Berlin, Heidelberg 2006), p. 228–235. Springer-Verlag.
- Rui Liao, Yunhao Tan, Hari Sundar, Marcus Pfister, and Ali Kamen. 2010. “An Efficient Graph-Based Deformable 2D/3D Registration Algorithm with Applications for Abdominal Aortic Aneurysm Interventions”. In *Medical Imaging and Augmented Reality*.
- Tony Lindeberg. 1996. “Edge Detection and Ridge Detection with Automatic Scale Selection”. In *CVPR*.
- J.D. Liu, D. Chelberg, C. Smith, and H. Chebrolu. 2007. “Distribution-based Level Set Segmentation for Brain MR Images”. In *BMVC07*.
- Jundong Liu. 2006. “Robust Image Segmentation using Local Median”. In *Proc. of the 3rd Canadian Conf. on Comp. Robot Vis.* p. 31.
- C. Lorenz, I. Carlsen, T. Buzug, C. Fassnacht, and J. Weese. 1997. Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images. *CVRMed-MRCAS’97*, volume 1205 of *LNCS*, p. 233–242. Springer Berlin / Heidelberg.
- L.M. Lorigo, O.D. Faugeras, W.E.L. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Wesin. 2001. “CURVES: Curve evolution for vessel segmentation”. *Medical Image Analysis*, vol. 5, p. 195–206.
- J.B.Antoine Maintz and Max A. Viergever. 1998. “A survey of medical image registration”. *Medical Image Analysis*, vol. 2, n. 1, p. 1–36.
- R. Malladi and J.A. Sethian. 1996. “A Unified Approach to Noise Removal, Image Enhancement, and Shape Recovery”. *IEEE IP*, vol. 5, p. 1554–1568.
- José V. Manjón, Juan J. Lull, José Carbonell-Caballero, Gracián García-Martí, Luís Martí-Bonmatí, and Montserrat Robles. 2007. “A nonparametric MRI inhomogeneity correction method”. *Medical Image Analysis*, vol. 11, n. 4, p. 336–345.
- R. Manniesing, B.K. Velthuis, M.S. van Leeuwen, I.C. van der Schaaf, P.J. van Laar, and W.J. Niessen. April 2006a. “Level set based cerebral vasculature segmentation and diameter quantification in CT angiography”. *Medical Image Analysis*, vol. 10, n. 2, p. 200–214.
- R. Manniesing, M. A. Viergever, and W. J. Niessen. Decembre 2006b. “Vessel enhancing diffusion: A scale space representation of vessel structures”. *Medical Image Analysis*, vol. 10, n. 6, p. 200–214.

- P. Markelj, D. Tomazevic, F. Pernus, and B. Likar. 2008. "Robust Gradient-Based 3-D/2-D Registration of CT and MR to X-Ray Images". *Medical Imaging, IEEE Transactions on*, vol. 27, n. 12, p. 1704–1714.
- P. Markelj, D. Tomazevic, B. Likar, and F. Pernus. April 2012. "A review of 3D/2D registration methods for image-guided interventions". *Medical Image Analysis*, vol. 16, n. 3, p. 642–661.
- José L. Marroquín, Baba C. Vemuri, Salvador Botello, and F. Calderon. 2002. "An Accurate and Efficient Bayesian Method for Automatic Segmentation of Brain MRI". *IEEE Trans. Med. Imag.*, vol. 21, p. 934-945.
- Graeme F. Mason and Douglas L. Rothman. 2002. "Graded image segmentation of brain tissue in the presence of inhomogeneous radio frequency fields". *Magnetic Resonance Imaging*, vol. 20, p. 431–436.
- Arnaldo Mayer and Hayit Greenspan. 2009. "An adaptive mean-shift framework for MRI brain segmentation". *IEEE TMI*, vol. 28, p. 1238–50.
- Erik H. W. Meijering, Wiro J. Niessen, Josien P. W. Pluim, and Max A. Viergever. 1999. "Quantitative Comparison of Sinc-Approximating Kernels for Medical Image Interpolation.". In *MICCAI'99*. p. 210-217.
- John Melonakos, Yi Gao, and Allen Tannenbaum. 2007. "Tissue tracking: applications for brain MRI classification". In *SPIE Med. Imag.*
- Coert Metz, Michiel Schaap, Stefan Klein, Lisan Neefjes, Ermanno Capuano, Carl Schultz, Robert Jan van Geuns, Patrick W. Serruys, Theo van Walsum, and Wiro J. Niessen. 2009. "Patient Specific 4D Coronary Models from ECG-gated CTA Data for Intra-operative Dynamic Alignment of CTA with X-ray Images". In *MICCAI 2009*. p. 369-376.
- B. Mory, R. Ardon, and J.-P. Thiran. 2007. "Variational Segmentation using Fuzzy Region Competition and Local Non-Parametric Probability Density Functions". In *ICCV*.
- David Mumford and J. Shah. 1989. "Optimal approximation by piecewise smooth functions and associated variational problems". *Communication on Pure and Applied Mathematics*, vol. 42, n. 4, p. 577-685.
- Delphine Nain, Anthony Yezzi, and Greg Turk. 2004. "Vessel Segmentation Using A Shape Driven Flow". In *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2004*.
- J. A. Nelder and R. Mead. 1965. "A simplex method for function minimization". *The Computer Journal*, vol. 7, n. 4, p. 308–313.
- Ca nero C. and Radeva P. 2003. "Vesselness enhancement diffusion". *Pattern Recognition Letters*, vol. 24, p. 3141–3151.

- H.J. Noordmans and A.W.M. Smeulders. 1998. "High accuracy tracking of 2D/3D curved line-structures by consecutive cross-section matching". *Pattern Recognition Letters*, vol. 19, n. 1, p. 97 - 111.
- B. Obara, M. Fricker, D. Gavaghan, and V. Grau. may 2012. "Contrast-Independent Curvilinear Structure Detection in Biomedical Images". *Image Processing, IEEE Transactions on*, vol. 21, n. 5, p. 2572–2581.
- Stanley Osher and J. A. Sethian. 1988. "Fronts Propagating with Curvature Dependent Speed: Algorithms Based on Hamilton–Jacobi Formulations". *Journal of Computational Physics*, vol. 79, n. 2, p. 12-49.
- Nikos Paragios, 2006. *Curve Propagation, Level Set Methods and Grouping, Handbook of Mathematical Models in Computer Vision*, chapter 9, p. 147-162. Springer.
- Nikos Paragios and Rachid Deriche. 2002. "Geodesic Active Regions: A new framework to deal with frame partition problems in Computer Vision". *J. Visual Commun. Image Rep.*, vol. 13, n. 2, p. 249-268.
- Nikos Paragios and Rachid Deriche, 2005. *Mathematical Models in Computer Vision: The Handbook*, chapter Level Set Methods & Visual Grouping, p. 147-162. Springer.
- Nikos Paragios, Mikael Rousson, and Visvanathan Ramesh. February 2003. "Non-rigid registration using distance functions". *Comput. Vis. Image Underst.*, vol. 89, p. 142–165.
- Ting Peng, Ian H. Jermyn, Véronique Prinet, and Josiane Zerubia. 2008. "Extended Phase Field Higher-Order Active Contour Model for Networks and its Application to Road Network Extraction from VHR Satellite Images". In *ECCV2008*.
- X. Pennec and J.-P. Thirion. 1997. "A framework for uncertainty and validation of 3-D registration methods based on points and frames". *Int. J. Comput. Vision*, vol. 25, n. 3, p. 203–229.
- Terry Peters and Kevin Cleary, 2008. *Image-Guided Interventions: Technology and Applications*. New York : Springer Science+Business Media.
- Terry M Peters. 2006. "Image-guidance for surgical procedures". *Physics in Medicine and Biology*, vol. 51, p. R505–R540.
- Dzung L. Pham and Jerry L. Prince. 1999. "Adaptive Fuzzy Segmentation of Magnetic Resonance Images". *IEEE TMI*, vol. 18, p. 737–752.
- Dzung L. Pham and Jerry L. Prince. 2004. "ROBUST UNSUPERVISED TISSUE CLASSIFICATION IN MR IMAGES". In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro, 2004*. p. 109–112.
- M. J. D. Powell. 1964. "An efficient method for finding the minimum of a function of several variables without calculating derivatives". *The computer journal*, vol. 7, n. 2, p. 155–162.

- M. J. D. Powell, 1994. *Advances in Optimization and Numerical Analysis*, chapter A direct search optimization method that models the objective and constraint functions by linear interpolation, p. 51–67. Kluwer Academic, Dordrecht.
- M. J. D. Powell. 1998. *Direct search algorithms for optimization calculations*. Technical report. Department of Applied Mathematics and Theoretical Physics, Cambridge University, UK.
- M. J. D. Powell. 2009. *The BOBYQA algorithm for bound constrained optimization without derivatives*. Technical report. Centre for Mathematical Sciences, Cambridge University, UK.
- M. Prümmer, J. Hornegger, M. Pfister, and A. Dörfler. 2006. “Multi-modal 2D-3D Non-rigid Registration”. In *SPIE Medical Imaging 2006*.
- Xiaoning Qian, Matthew P. Brennan, Donald P. Dione, Wawrzyniec L. Dobrucki, Marcel P. Jackowski, Christopher K. Breuer, Albert J. Sinusas, and Xenophon Papademetris. 2009. “A non-parametric vessel detection method for complex vascular structures”. *Medical Image Analysis*, vol. 13, p. 49–61.
- Jagath C. Rajapakse and Frithjof Kruggel. 1998. “Segmentation of MR images with intensity inhomogeneities”. *Image and Vision Computing*, vol. 16, n. 3, p. 165–180.
- D. Rivest-Hénault, L. Duong, C. Lapierre, S. Deschênes, and M. Cheriet. 2010a. “Semi-automatic segmentation of major aorto-pulmonary collateral arteries for image guided procedures”. In *SPIE Med. Imag.*
- D. Rivest-Hénault, C. Lapierre, S. Deschênes, and M. Cheriet. 2010b. “Length increasing active contour for the segmentation of small blood vessels”. In *20th International Conference on Pattern Recognition*.
- David Rivest-Hénault and Mohamed Cheriet. 2007. “Image Segmentation Using Level Set and Local Linear Approximations”. In *ICIAR 2007*. p. 234-245. Springer-Verlag.
- David Rivest-Hénault and Mohamed Cheriet. 2012. “Vessel detection filter via structure-ball analysis”. *Submitted to IEEE Image Processing*, vol. 0, p. 0.
- David Rivest-Hénault, Reza Farrahi Moghaddam, and Mohamed Cheriet. 2011. “A local linear level set method for the binarization of degraded historical document images”. *International Journal on Document Analysis and Recognition*, vol. Online First, p. 1–24.
- David Rivest-Hénault, Hari Sundar, and Mohamed Cheriet. 2012. “Non-Rigid 2D/3D Registration of Coronary Artery Models with Live Fluoroscopy for Guidance of Cardiac Interventions”. *IEEE-TMI*, vol. IN PRESS, p. -.
- Marie Rochery, Ian Jermyn, and Josiane Zerubia. 2005. “Phase Field Models and Higher-Order Active Contours”. In *ICCV2005*.

- Wayne Rosamond, Katherine Flegal, Karen Furie, Alan Go, Kurt Greenlund, Nancy Haase, Susan M. Hailpern, Michael Ho, Virginia Howard, Bret Kissela, Steven Kittner, Donald Lloyd-Jones, Mary McDermott, James Meigs, Claudia Moy, Graham Nichol, Christopher O'Donnell, Veronique Roger, Paul Sorlie, Julia Steinberger, Thomas Thom, Matt Wilson, Yuling Hong, for the American Heart Association Statistics Committee, and Stroke Statistics Subcommittee. 2008. "Heart Disease and Stroke Statistics—2008 Update: A Report From the American Heart Association Statistics Committee and Stroke Statistics Subcommittee". *Circulation*, vol. 117, n. 4, p. e25-146.
- Bodo Rosenhahn, Thomas Brox, and Joachim Weickert. 2007. "Three-Dimensional Shape Knowledge for Joint Image Segmentation and Pose Tracking". *Int J Comput Vis*, vol. 73, n. 3, p. 243–262.
- T. Rowan. 1990. "Functional Stability Analysis of Numerical Algorithms". PhD thesis, Department of Computer Sciences, University of Texas at Austin.
- S. Roy and I. Cox. 1998. "A maximum-flow formulation of the n-camera stereo correspondence problem.". In *IEEE Proc. of Int. Conference on Computer Vision*.
- Daniel Ruijters, Bart M. ter Haar Romeny, and Paul Suetens. 2009. "Vesselness-based 2D–3D registration of the coronary arteries". *Int J CARS*, vol. 4, p. 391–397.
- Tomoko Sakai, Akichika Mikami, Masaki Tomonaga, Mie Matsui, Juri Suzuki, Yuzuru Hamada, Masayuki Tanaka, Takako Miyabe-Nishiwaki, Haruyuki Makishima, Masato Nakatsukasa, and Tetsuro Matsuzawa. 2011. "Differential Prefrontal White Matter Development in Chimpanzees and Humans". *Current Biology*, vol. 21, n. 16, p. 1397 - 1402.
- Yoshinobu Sato, Shin Nakajima, Nobuyuki Shiraga, Hideki Atsumi, Shigeyuki Yoshida, Thomas Koller, Guido Gerig, and Ron Kikinis. 1998. "Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images". *Medical Image Analysis*, vol. 2, n. 2, p. 143 - 168.
- Michiel Schaap, Coert T. Metz, Theo van Walsum, Alina G. van der Giessen, Annick C. Weustink, Nico R. Mollet, Christian Bauer, Hrvoje Bogunović, Carlos Castro, Xiang Deng, Engin Dikici, Thomas O'Donnell, Michel Frenay, Ola Friman, Marcela Hernández Hoyos, Pieter H. Kitslaar, Karl Krissian, Caroline Kühnel, Miguel A. Luengo-Oroz, Maciej Orkisz, Örjan Smedby, Martin Styner, Andrzej Szymczak, Hüseyin Tek, Chunliang Wang, Simon K. Warfield, Sebastian Zambal, Yong Zhang, Gabriel P. Krestin, and Wiro J. Niessen. 2009a. "Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms". *Medical Image Analysis*, vol. 13, n. 5, p. 701–714.
- Michiel Schaap, Lisan Neefjes, Coert Metz, Alina van der Giessen, Annick Weustink, Nico Mollet, Jolanda Wentzel, Theo van Walsum, and Wiro Niessen. 2009b. Coronary lumen segmentation using graph cuts and robust kernel regression. Prince, J., Dzung Pham, and Kyle Myers, editors, *Information Processing in Medical Imaging*, volume 5636 of

Lecture Notes in Computer Science, p. 528–539. Springer Berlin / Heidelberg. ISBN 978-3-642-02497-9.

Matthias Schneider and Hari Sundar. 2010. “Automatic global vessel segmentation and catheter removal using local geometry information and vector field integration”. In *IEEE ISBI 2010*.

J. Sethian. 1996. “A Marching Level Set Method for Monotonically Advancing Fronts”. In *Proceedings of the National Academy of Science*.

J.A. Sethian, 1999. *Level Set Methods and Fast Marching Methods*. Cambridge University Press.

Ken Shoemake and Tom Duff. 1992. “Matrix Animation and Polar Decomposition”. In *Proceedings of the conference on Graphics interface '92*. p. 258–264. Morgan Kaufmann Publishers Inc.

Karan Sikka, Nitesh Sinha, Pankaj K. Singh, and Amit K. Mishra. 2009. “A fully automated algorithm under modified FCM framework for improved brain MR image segmentation”. *Magnetic Resonance Imaging*, vol. 27, p. 994–1004.

M. Y. Siyal and Lin Yu. October 2005. “An intelligent modified fuzzy c-means based algorithm for bias estimation and segmentation of brain MRI”. *Pattern Recognition Letters*, vol. 26, n. 13, p. 2052–2062.

John G. Sled, Alex P. Zijdenbos, and Alan C. Evans. Feb 1998. “A nonparametric method for automatic correction of intensity nonuniformity in MRI data”. *IEEE Tran. Med. Imag.*, vol. 17, n. 1, p. 87–97.

Peter-Pike Sloan. 2008. “Stupid Spherical Harmonics (SH) Tricks”. Game Developers Conference 2008.

Laurent Sarry Souha Aouadi. 2007. “Accurate and precise 2D–3D registration based on X-ray intensity”. *Computer Vision and Image Understanding*, vol. 110, p. 134–151.

C. Stewart, 2006. *Handbook of Mathematical Models in Computer Vision*, chapter 14 — Uncertainty-Driven, Point-Based Image Registration, p. 221–235. Springer.

Gregg W. Stone, Nicolaus J. Reifart, Issam Moussa, Angela Hoyer, David A. Cox, Antonio Colombo, Donald S. Baim, Paul S. Teirstein, Bradley H. Strauss, Matthew Selmon, Gary S. Mintz, Osamu Katoh, Kazuaki Mitsudo, Takahiko Suzuki, Hideo Tamai, Eberhard Grube, Louis A. Cannon, David E. Kandzari, Mark Reisman, Robert S. Schwartz, Steven Bailey, George Dangas, Roxana Mehran, Alexander Abizaid, Jeffrey W. Moses, Martin B. Leon, and Patrick W. Serruys. 2005. “Percutaneous Recanalization of Chronically Occluded Coronary Arteries”. *Circulation*, vol. 112, n. 16, p. 2530–2537.

R. Storn and K. Price. 1997. “Differential Evolution – a Simple and Efficient Heuristic for Global Optimization over Continuous Spaces”. *Journal of Global Optimization*, vol. 11, p. 341–359.

- Hari Sundar, Ali Khamene, Chenyang Xu, Frank Sauer, and Christos Davatzikos. 2006. "A novel 2D-3D registration algorithm for aligning fluoro images with 3D pre-op CT/MR images". In *SPIE Med. Imag. 2006*.
- G. Sundaramoorthi, A. Yezzi, A. C. Mennucci, and G. Sapiro. 2009. "New Possibilities with Sobolev Active Contours". *Int. J. Comput. Vis.*, vol. 84, p. 113–129.
- Ganesh Sundaramoorthi and Anthony Yezzi. March 2007. "Global Regularizing Flows With Topology Preservation for Active Contours and Polygons". *IEEE Trans. Image Process.*, vol. 16, n. 3, p. 803–812.
- D. Tomazevic, B. Likar, and F. Pernus. jan. 2006. "3-D/2-D registration by integrating 2-D information in 3-D". *Medical Imaging, IEEE Transactions on*, vol. 25, n. 1, p. 17–27.
- Michael VN Truong, Abdullah Aslam, C Aldo Rinaldi, Reza Razavi, Graeme P Penney, and Kawal S Rhode. 2009. "Preliminary Investigation: 2D–3D Registration of MR and X-ray Cardiac Images Using Catheter Constraints". In *CI2BM09 - MICCAI Workshop on Cardiovascular Interventional Imaging and Biophysical Modelling (2009)*.
- Andy Tsai, Anthony Yezzi, and Alan S. Willsky. August 2001. "Curve Evolution Implementation of the Mumford–Shah Functional for Image Segmentation, Denoising, Interpolation, and Magnification". *IEEE Trans. Image Process.*, vol. Volume 10, n. 8, p. 1169–1186.
- Yanghai Tsin, Klaus Kirchberg, Guenter Lauritsch, and Chenyang Xu. 2009. "A Deformation Tracking Approach to 4D Coronary Artery Tree Reconstruction". In *MICCAI 2009*.
- David S. Tuch. 2004. "Q-ball imaging". *Magnetic Resonance in Medicine*, vol. 52, n. 6, p. 1358–1372.
- Guy-Anne Turgeon, Glen Lehmann, Gerard Guiraudon, Maria Drangova, David Holdsworth, and Terry Peters. 2005. "2D-3D registration of coronary angiograms for cardiac procedure planning and guidance". *Med. Phys.*, vol. 32, p. 3737–3749.
- M. Vaidyanathan, L. P. Clarke, C. Heidtman, R. P. Velthuisen, and L. O. Hall. 1997. "Normal Brain Volume Measurement Using Multispectral MRI Segmentation". *Magnetic Resonance Imaging*, vol. 15, p. 87–97.
- C. M. van Bommel, L. J. Spreeuwiers, M. A. Viergever, and W. J. Niessen. 2003. "Level-Set-Based Artery–Vein Separation in Blood Pool Agent CE-MR Angiograms". *IEEE Trans. Med. Imag.*, vol. 22, p. 1224–1234.
- I. M. J. van der Bom, S. Klein, M. Staring, R. Homan, L. W. Bartels, and J. P. W. Pluim. 2011. "Evaluation of optimization methods for intensity-based 2D–3D registration in X-ray guided interventions". In *SPIE Medical Imaging*.
- K. van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. 1999. "Automated model-based tissue classification of MR images of the brain". *IEEE Trans. Med. Imag.*, vol. 18, p. 897–908.

- Alexander Vasilevskiy and Kaleem Siddiqi. 2002. "Flux Maximizing Geometric Flows". *IEEE PAMI*, vol. 24, n. 12, p. 1565–1578.
- Annelena Venneri. 2007. "Imaging treatment effects in ALzheimer's disease". *Magnetic Resonance Imaging*, vol. 25, p. 953–968.
- Luminita Vese, 2003. *Multiphase Object Detection and Image Segmentation, Geometric Level Set Methods*, chapter 10, p. 175-194. Springer.
- Luminita A. Vese and Tony F. Chan. 2002. "A multiphase Level Set Framework for Image Segmentation Using the Mumford and Shah Model". *International Journal of Computer Vision*, vol. Volume 50, n. 2, p. 271–293.
- Mette R. Wiegell, Henrik B. W. Larsson, and Van J. Wedeen. 2000. "Fiber Crossing in Human Brain Depicted with Diffusion Tensor MR Imaging". *Radiology*, vol. 217, p. 897–903.
- W.C.K. Wong, R.W.K. So, and A.C.S. Chung. april 2012. "Principal Curves for Lumen Center Extraction and Flow Channel Width Estimation in 3-D Arterial Networks: Theory, Algorithm, and Validation". *Image Processing, IEEE Transactions on*, vol. 21, n. 4, p. 1847–1862.
- World Health Organization. 2008. "Cardiovascular diseases". <http://www.who.int/cardiovascular_diseases/en/>. accessed 2-April-2008.
- A. Worth. 1996. "Internet Brain Segmentation Repository". <<http://www.cma.mgh.harvard.edu/ibsr/>>.
- S.. Wörz and K.. Rohr. aug. 2007. "Segmentation and Quantification of Human Vessels Using a 3-D Cylindrical Intensity Model". *Image Processing, IEEE Transactions on*, vol. 16, n. 8, p. 1994 -2004.
- Stefan Wörz, William J. Godinez, and Karl Rohr. 2009. Probabilistic tracking and model-based segmentation of 3d tubular structures. *Bildverarbeitung für die Medizin 2009, Informatik aktuell*, p. 41–45. Springer Berlin Heidelberg.
- Yaddah. 2006. "Diagram of the human heart (cropped)". On line. <http://en.wikipedia.org/wiki/File:Diagram_of_the_human_heart_%28cropped%29.svg>. Published under a permissive Creative Common License.
- Yuan Yuan, Yishan Luo, and A.C.S. Chung. july 2011. "VE-LLI-VO: Vessel Enhancement Using Local Line Integrals and Variational Optimization". *Image Processing, IEEE Transactions on*, vol. 20, n. 7, p. 1912 -1924.
- Yingjie Zhang. 2005. "Fast Segmentation for the Piecewise Smooth Mumford-Shah Functional". *International Journal of Signal Processing*, vol. 2, p. 245-250.
- Yingjie Zhang. 2006. "An Improved Algorithm for the Piecewise-Smooth Mumford and Shah Model in Image Segmentation". *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2006, p. 1-5.

- Yongyue Zhang, Michael Brady, and Stephen Smith. 2001. "Segmentation of Brain MR Images Through a Hidden Markov Random Field Model and the Expectation-Maximization Algorithm". *IEEE TMI*, vol. 20, p. 45–57.
- Fei Zhao, Honghai Zhang, Andreas Wahle, Matthew T. Thomas, Alan H. Stolpen, Thomas D. Scholz, and Milan Sonka. 2009. "Congenital aortic disease: 4D magnetic resonance segmentation and quantitative analysis". *Medical Image Analysis*, vol. 13, p. 483–493.
- Guoyan Zheng, Miguel Á.G. Ballester, Martin Styner, and Lutz-Peter Nolte. 2006. Reconstruction of patient-specific 3d bone surface from 2d calibrated fluoroscopic images and point distribution model. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, volume 4190 of *Lecture Notes in Computer Science*, p. 25-32. Springer Berlin Heidelberg.
- Guoyan Zheng, Sebastian Gollmer, Steffen Schumann, Xiao Dong, Thomas Feilkas, and Miguel A. González Ballester. 2009. "A 2D/3D correspondence building method for reconstruction of a patient-specific 3D bone surface model using point distribution models and calibrated X-ray images". *Medical Image Analysis*, vol. 13, n. 6, p. 883 - 899.
- Yefeng Zheng, Maciej Loziczzonek, Bogdan Georgescu, S. Kevin Zhou, Fernando Vega-Higuera, and Dorin Comaniciu. 2011. "Machine learning based vesselness measurement for coronary artery segmentation in cardiac CT volumes". In *SPIE Medical Imaging*. p. 79621K. SPIE.
- Zhenyu Zhou and Zongcai Ruan. 2007. "Multicontext wavelet-based thresholding segmentation of brain tissues in magnetic resonance images". *Magnetic Renonance Imaging*, vol. 25, p. 381–385.