
Myocardial Motion and Deformation Analysis from Echocardiograms

Michael Sühling

Thèse N° 3049 (2004)

*Thèse présentée à la faculté des sciences et techniques de l'ingénieur
pour l'obtention du grade de docteur ès sciences
et acceptée sur proposition du jury*

*Prof. René Salathé, président
Prof. Michael Unser, directeur de thèse
Prof. Aude Billard, rapporteur
Dr Patrick Hunziker, rapporteur
Prof. Rudolf Mester, rapporteur*

École polytechnique fédérale de Lausanne—2004

*Cover design by Annette Unser
Printing and binding by Repro-EPFL
Typeset with L^AT_EX
Copyright © 2004 by Michael Sühling
Available at <http://bigwww.epfl.ch/>*

Abstract

Echocardiography is a widely used imaging technique to examine myocardial function in patients with known or suspected heart disease. The analysis of ventricular wall motion and deformation, in particular, allows to assess the extent of myocardial ischemia and infarction. In clinical practice, the analysis mainly relies on visual inspection or manual measurements by experienced cardiologists. Manual methods are tedious and time-consuming, and visual assessment leads to qualitative and subjective diagnoses that suffer from considerable inter- and intraobserver variability. Automating the analysis of echocardiographic images is therefore highly desirable but also challenging because of the low image quality and the high amount of speckle noise. In this thesis, we propose a framework for robust and quantitative analysis of echocardiographic sequences. We make the following key contributions:

Motion and Deformation Analysis—We propose a novel optical-flow-based algorithm to estimate ventricular wall motion from B-mode echocardiograms. To account for typical heart motions such as contraction/expansion and shear, we use a local affine model for the velocity in space and time. An attractive feature of the affine motion model is that it gives also access to local strain rate parameters that describe local myocardial deformation such as wall thickening. The motion parameters are estimated in the least-squares sense within a sliding spatio-temporal B-spline window. The estimation of large motions is made possible through the use of a coarse-to-fine multi-scale strategy, which also adds robustness to the method.

Computational Efficiency—We introduce the notion of multiresolution moment filters, a novel filtering scheme to compute local weighted geometric moments efficiently at dyadic scales by using a wavelet-like algorithm. Beyond their application in motion analysis, we demonstrate their usefulness for image denoising and feature extraction.

Multi-Modality—We extend the proposed motion analysis algorithm by inte-

grating directional, Doppler-based velocity measurements. The exploitation of two ultrasound modalities, i.e., B-mode and tissue Doppler, renders the method more accurate and robust.

Visualization—We display diagnostically meaningful motion data inside a user-defined region of interest that is tracked in time. Myocardial inward and outward motion is visualized by color coding the radial motion component with respect to the ventricular center. Two-dimensional strain rate information is superimposed in the form of deforming ellipses. The display allows a more intuitive and simplified identification of regions with abnormal motion patterns.

Validation—The proposed method is validated on 1) synthetic data, 2) real ultrasound phantom data, and 3) clinical echocardiograms. A large-scale validation study that includes 114 patients confirms its ability to detect and quantify wall motion abnormalities.

Version Abrégée

L'échocardiographie est une technique d'imagerie qui est très souvent utilisée pour examiner le fonctionnement du myocarde de patients potentiellement sujets à une maladie cardiaque. En particulier, l'analyse du mouvement et de la déformation de la paroi du ventricule permet d'évaluer le degré d'ischémie et d'infarctus du myocarde. Dans la pratique clinique, cette analyse repose principalement sur l'inspection visuelle et sur les mesures manuelles réalisées par des cardiologues expérimentés. Or, les méthodes manuelles sont longues et pénibles; quant à l'évaluation visuelle, elle conduit à des diagnostics qualitatifs et subjectifs qui souffrent considérablement de la variabilité inter et intra-observateur. Automatiser l'analyse des images par échocardiographie est donc fortement souhaitable, mais aussi difficile en raison de la médiocre qualité des images et de la grande quantité de bruit de granularité. Dans cette thèse, nous proposons un ensemble de méthodes pour analyser de façon robuste et quantitative les séquences d'image en échocardiographie. Nos contributions principales sont les suivantes:

Analyse des mouvements et des déformations—Nous proposons un algorithme basé sur le flot optique pour estimer le mouvement de la paroi du ventricule à partir d'échocardiogrammes B-mode. Pour bien tenir compte de certains mouvements typiques du cœur tels que contraction/expansion et cisaillement, nous décrivons la vitesse dans l'espace et le temps à l'aide d'un modèle affine local. Un aspect intéressant du modèle de mouvement affine est qu'il permet aussi de connaître localement les contraintes associées à certaines déformations locales du myocarde, telles que l'épaississement de sa paroi. Au sein d'une fenêtre B-spline spatio-temporelle glissante, nous estimons les paramètres de mouvement au sens des moindres carrés. L'utilisation d'une stratégie multi-échelle grossière-à-fine tout à la fois rend la méthode plus robuste et nous permet d'évaluer les mouvements de grande amplitude.

Efficacité de calcul—Nous présentons la notion de filtre multi-échelle de mo-

ments, qui consiste en une nouvelle structure de filtres à échelle dyadique destinés au calcul efficace de moments géométriques avec pondération locale, et qui emploie un algorithme similaire à celui des ondelettes. Outre leur application à l'analyse de mouvement, nous démontrons leur utilité pour le débruitage d'image et pour l'extraction de caractéristiques.

Multi-modalité—Nous complétons l'algorithme d'analyse de mouvements proposé par des mesures de vitesse directionnelle basées sur l'effet Doppler. Nous rendons la méthode encore plus précise et plus robuste en exploitant simultanément les deux modalités d'ultrason que sont B-mode et Doppler tissulaire.

Visualisation—Nous affichons les données de mouvement pertinentes au diagnostic à l'intérieur d'une région d'intérêt qui est définie par l'utilisateur et qui est continûment mise à jour. Nous visualisons les mouvements centripètes et centrifuges du myocarde par un codage en fausses couleurs de la composante radiale du mouvement par rapport au centre du ventricule. La représentation bidimensionnelle des contraintes s'y superpose sous la forme d'ellipses déformables. Cet affichage permet une identification plus facile et plus intuitive des régions qui comportent des mouvements anormaux.

Validation—Nous validons la méthode proposée au moyen 1) de données synthétiques, 2) de données réelles d'ultrasons de fantômes et 3) d'échocardiogrammes cliniques. Une étude de validation à large échelle, qui inclut 114 patients, confirme la capacité de notre méthode à détecter et à quantifier les anomalies de mouvement de la paroi du ventricule.

To my parents.

Acknowledgement

This thesis describes the results of the research carried out during my Ph.D. studies at the Biomedical Imaging Group at the Swiss Federal Institute of Technology Lausanne. Certainly, this work would not have been possible without the support of many people.

First of all, I would like to express my gratitude to my thesis advisor Prof. Michael Unser. The discussions with him and his constructive criticism in writing papers were very crucial. Most of all, I appreciate his enthusiasm, creativity and global view in conducting research; I hope that I could internalize at least parts of it.

Many thanks also go to Dr. Patrick Hunziker from the University Hospital Basel who initiated this interesting research project¹. His overwhelming source of ideas was extremely stimulating. The close collaboration with him allowed me to gain a deep insight into the clinical background and the various problems originating from it.

I would like to thank all current and former members of the Biomedical Imaging Group for the pleasant working environment; I greatly appreciated to interact with all of them. Special thanks go to Muthuvel Arigovindan with whom I shared my office during all these years.

I also gratefully acknowledge all people with whom I cooperated during my numerous stays at the University Hospital Basel. Specifically, I would like to thank Dr. Christian Jansen with whom I worked very closely and without him the clinical validation experiments would not have been possible.

¹The project was partially funded by the Swiss Heart Foundation and the Swiss National Science Foundation under grant 3200-059517.99.

Last but not least, I would like to thank my whole family, in particular, my parents; they have always supported and encouraged me in my work, which allowed me to get to the point where I am.

Michael Sühling
Lausanne, July 2004

Contents

Abstract	i
Version Abrégée	iii
Acknowledgement	vii
1 Introduction	1
1.1 Cardiac Motion Patterns	2
1.2 Related Work	3
1.3 Thesis Organization	5
2 Basic Heart Anatomy and Physiology	9
2.1 Introduction	9
2.2 Heart Anatomy and Normal Functioning	10
2.2.1 Heart Anatomy	10
2.2.2 Normal Physiology	12
2.2.3 Electrophysiology of the Heart	13
2.2.4 Coronary Circulation	14
2.3 Common Cardiovascular Diseases	15
2.3.1 Coronary Artery Disease	15
2.3.2 Other Diseases	16
2.4 Conventional Diagnostic Assessment of Myocardial Malfunction	17
3 Principles and Instrumentation of Echocardiography	19
3.1 Introduction	19
3.2 Basic Principles of Ultrasound	21
3.2.1 Imaging Modes	23

3.2.2	Doppler Echocardiography	25
3.3	Standard Views for Two-Dimensional Echocardiography	27
3.3.1	Apical Views	28
3.3.2	Parasternal Views	29
3.3.3	Echocardiographic Nomenclature	29
3.4	Conclusions	33
4	Motion Analysis from B-mode Echocardiograms	35
4.1	Introduction	36
4.2	Local-Affine, Multi-Scale Motion Estimation	37
4.2.1	Local Affine Velocity in Space-Time	38
4.2.2	Choice of Window Size—Multiresolution Moments Computation	39
4.2.3	Coarse-To-Fine Multi-Scale Strategy	43
4.3	ROI-Tracking and Motion Visualization	45
4.4	Numerical Results	47
4.4.1	Application to Synthetic Data	48
4.4.2	Phantom Experiments	54
4.4.3	Application to Clinical Data	57
4.5	Conclusions	63
5	Total-Least-Squares-Based Motion Estimation	65
5.1	Introduction	65
5.2	The Total-Least-Squares Problem	67
5.2.1	Solution of the Total-Least-Squares Problem	68
5.2.2	Relation between Least-Squares and Total-Least-Squares	69
5.2.3	Scaled Total-Least-Squares	70
5.2.4	Equilibration of the Total-Least-Squares System	71
5.3	Numerical Results	72
5.4	Conclusions	75
6	Bimodal Motion Analysis from B-mode and Tissue Doppler Ultrasound	77
6.1	Introduction	78
6.2	Principles of Tissue Doppler Imaging	79
6.3	Motion Analysis from B-mode Echocardiograms	79
6.4	Bimodal Motion Analysis	80
6.4.1	Coarse-To-Fine Multi-Scale Strategy	82
6.5	Numerical Results	83

6.5.1	Application to Synthetic Data	83
6.5.2	Phantom Experiments	84
6.5.3	Application to Clinical Data	84
6.6	Conclusions	86
7	Strain Rate Analysis	87
7.1	Introduction	87
7.2	Definition of Strain and Strain Rate	88
7.3	Two-Dimensional Strain Rate Analysis	90
7.3.1	Incompressibility Assumption	92
7.3.2	Strain Rate Visualization	92
7.4	Numerical Results	93
7.4.1	Application to Synthetic Data	93
7.4.2	Application to Clinical Data	97
7.5	Conclusions	98
8	Clinical Validation	99
8.1	Introduction	99
8.2	Experimental Setup and Statistical Validation	100
8.2.1	Comparison with Tissue Doppler	102
8.2.2	Segmental Multiple Comparison Tests	103
8.2.3	Artificial Neural Network Classification	108
8.3	Exemplary Case Studies	110
8.3.1	Comprehensive Objective Analysis of Paradoxical Motion	111
8.3.2	New Insights into Cardiac Motion	111
8.4	Conclusions	113
9	Conclusion	115
9.1	Main Contributions	115
9.2	Future Research Directions	117
A	Multiresolution Moment Filters	119
A.1	Introduction	120
A.2	Theory	121
A.2.1	Weighted Local Geometric Moments	121
A.2.2	Two-Scale Equation	122
A.2.3	Efficient Multi-Scale Implementation	123
A.2.4	Multiple Dimensions	126
A.2.5	B-spline Window Function	127

A.3	Applications	129
A.3.1	Local Shape Analysis and Feature Extraction	129
A.3.2	Multi-Scale, Weighted Savitzky-Golay Smoothing Filters .	135
A.3.3	Optical-Flow Estimation	140
A.4	Conclusions	145
A.5	Appendix	145
A.5.1	Computational Complexity	145
A.5.2	Computation of the Probability Density Function of a Weighted Least-Squares Residual	147
	Bibliography	149
	Curriculum Vitæ	161

A good heart is worth gold.

— William Shakespeare

Chapter 1

Introduction

The heart is the most active organ in the human body. Beating about every second throughout our lifetime, it continuously supplies the body with vital blood. The heart has always played a key role in medical science and has even become a major symbol for life. Its life-sustaining pumping function is realized by its powerful muscular structure—the so-called *myocardium*. Its activity is kept up by an efficient circulatory system of coronary arteries that supplies the muscle with the essential oxygen and nutrients.

Silent but progressive occlusion and hardening of the coronary arteries is the main cause for cardiac malfunction. It is caused by a build-up of fatty deposits in the lining of the artery walls—a process known as *atherosclerosis*. The lack of blood supply—the so-called *ischemia*—reduces the heart’s ability to contract normally. When one or more of the coronary arteries get completely occluded, blood to the heart muscle is cut off and the affected myocardial segments get seriously damaged or even die. This so-called *myocardial infarction*, also known as “heart attack”, often occurs suddenly and may be life-threatening. Coronary artery disease—specifically, heart attack—is the main cause of long-term disability and death throughout the industrial world. In Europe, about four million people die of cardiovascular disease every year, which is claiming more lives than the next five leading causes of death altogether.

Modern technology provides a number of tools to access myocardial function. Echocardiography, in particular, has evolved as a well-established imaging technique for the non-invasive evaluation of regional and global wall motion

abnormalities. Whereas echocardiographic imaging has reached a high level of technical sophistication, the analysis of ventricular function itself still relies mainly on the visual assessment and manual measurements by experienced cardiologists. According to a standardization of the American Heart Association, the function of myocardial segments is usually assigned to one of five descriptive categories, ranging from hyperactive contraction via virtually absent activity to paradoxical, counterproductive wall motion. The presence of segmental variations in contractility adds an additional level of difficulty to this semi-quantitative analysis. Consequently, the diagnosis depends highly on the observer's subjective interpretation and thus suffers from a considerable degree of inter- and intraobserver variability.

The ability to assess heart function in an objective and reproducible manner is therefore highly needed for a number of reasons: the detection and quantification of impaired myocardial function and the ability to objectively measure the impact of therapeutic treatments may considerably increase the efficiency of clinical interventions. Measuring ventricular contraction at different levels of physical stress may render non-invasive diagnostic methods such as stress echocardiography more effective, allowing the detection of myocardial malfunction at an earlier stage.

Since the left ventricle provides the whole body with oxygen-rich blood, it is the strongest and most dominant part of the heart. Therefore, during the analysis of heart motion, most of the attention is drawn to this chamber whose role is absolutely vital.

1.1 Cardiac Motion Patterns

The nature of myocardial function is complex and encompasses various motion patterns. Conceptually, the two most basic aspects of myocardial function are: wall *displacement* and *deformation*. Fig. 1.1 shows a typical echocardiogram of the left ventricle during contraction. As indicated in Fig. 1.1a, the myocardium contracts, moving inwards to eject the blood into the body. Measuring the velocity of the inward motion yields valuable information to characterize ventricular function.

In addition to the global inward motion, the myocardium usually thickens considerably to further increase its pumping efficiency. This deformation is highlighted in Fig. 1.1b. Additional deformations also include a longitudinal and circumferential shortening of the muscle fibers. The rate of deformation may be an indicator of regional contractility and is reflected by the velocity




Figure 1.1: Basic motion patterns of contracting ventricle. The myocardium moves inwards (a) and thickens significantly (b) to maximize pumping efficiency.

gradient between neighboring tissue segments. The quantitative assessment of regional deformation rates is likely to be the key for distinguishing between actively contracting segments and infarcted ones which are completely passive.

The most successful current technique—tissue Doppler-based analysis of motion and deformation—is hampered by the critical dependence on the ultrasound beam direction, leading to signal dropouts and velocity underestimation. To date, there has been a considerable amount of research on the quantification of heart motion from ultrasound data, but there is not yet a fully satisfactory solution. Some popular approaches are reviewed in the next section.

1.2 Related Work

There is a large literature describing the application of so-called *active contour* or *snake* techniques [1] to cardiac ultrasound [2, 3, 4]. These methods track the inner (endocardial) and/or outer (epicardial) border of the myocardium by minimizing a combination of internal and external energies of the contour representation. The internal energy is defined to keep the contour smooth, while

the external energy is designed to attract the contour towards the boundary. The external energy is usually composed of low-level image features such as high-intensity gradients. Since ultrasound images are typically noisy and of low quality, the inner and, especially, outer borders are usually not clearly defined, limiting the applicability of standard border detection algorithms significantly. The presence of additional structures such as the papillary muscles and the mitral valve creates a multitude of perturbing edges even in good quality images. Therefore, human guidance is often needed to guarantee acceptable results.

Active shape models (ASM) [5] are a generalization of active contours and were developed to overcome the latter problems. The key difference is that an active shape model can only undergo shape deformations that are consistent with a statistical model derived from some training data. Probabilistic shape-spaces are usually obtained using point distribution models (PDM) [6]. These models are determined by performing a principal component analysis (PCA) on landmark coordinates that are placed manually on many example images. The principal component analysis determines the strongest modes of shape variation which are then used as *a priori* information to constrain the contour deformation in sequences to be analyzed. Active shape models have been applied in echocardiography with partial success [7, 8, 9]. The major shortcoming of these border tracing approaches is that they only yield motion information normal to the myocardial boundaries. Motion parallel to the contours, such as longitudinal lengthening and shortening, and deformation of myocardial structures in-between the boundaries cannot be assessed.

The approach of active shape models was extended to the concept of active appearance models (AAM) by Cootes *et al.* [10]. Active appearance models describe both the image appearance, i.e., its gray value information, and shape as a combined statistical shape-appearance model. The approach was applied to cardiac motion analysis by [11]. Trainable-model-based approaches are extremely labour intensive because they require the manual tracing of a large set of echocardiograms. The training set should also include a sufficient number of cases with abnormal motion to be applicable for detecting myocardial malfunction. Furthermore, echocardiograms need to be normalized with respect to image size and number of frames using interpolation methods. Since the image brightness and contrast depend highly on the patient, the examiner and the ultrasound system used, the gray values of the echocardiograms need also to be normalized. The practical value of these approaches remains uncertain due to the limited testing on a wide range of clinical data.

Other motion analysis approaches are based on Markov-random-fields to segment cardiac ultrasound images [12, 13]. These methods take into account

ultrasound-specific noise models such as Rayleigh statistics. Since, however, the noise characteristics of commercial ultrasound systems are usually difficult to assess, the applicability of these methods to routine ultrasound remains limited.

Some researchers proposed to compute dense cardiac displacement fields using non-rigid registration techniques. Although elastic registration is applied in several areas of medical image analysis, its application to ultrasound is still limited [14, 15].

A related approach to cardiac motion assessment is based on the so-called *optical-flow* principle [16]. This technique estimates a dense velocity field of the moving image content from spatio-temporal intensity gradients. The instantaneous motion vectors are determined such that they point into the direction of minimum gray value changes. Application examples of this approach to echocardiograms yielded promising results [17, 18]. A popular optical-flow algorithm is the Lucas-Kanade method [19], which estimates the motion locally, assuming that the velocity field is constant within an observation window. This method was applied to ultrasound in [20]. An evaluation of different optical-flow methods applied to echocardiograms can be found in [21].

1.3 Thesis Organization

In this thesis, we propose a comprehensive approach to quantify different aspects of cardiac motion from echocardiograms independently from borders and Doppler. The thesis is organized as follows.

Following this general introduction, the thesis proceeds in Chapter 2 with an overview of the heart anatomy and of its normal function. We introduce technical terms that are frequently used throughout this work. We then describe symptoms and origins of typical wall motion abnormalities and review their conventional diagnostic assessment that is based on a semi-quantitative wall motion scoring.

In Chapter 3, we review the basic principles and instrumentation techniques of echocardiography. We give an overview of different imaging modes and introduce the standard two-dimensional views that are usually acquired during clinical examinations. We also introduce the echocardiographic nomenclature that is used to describe the different myocardial segments imaged.

The subsequent chapters present the scientific contributions of this work. Beginning in Chapter 4, we present a novel optical-flow-based method for estimating heart motion from two-dimensional B-mode echocardiograms. This sliding-window-type method uses a local affine velocity model that is partic-

ularly well suited to the description of typical heart motion. We employ a multiresolution strategy to be able to estimate large motions and to increase the robustness to noise. Motion models within windows of dyadic sizes are computed efficiently by using the concept of multiresolution moment filters—a novel, versatile filtering scheme that is discussed in detail in Appendix A at the end of this thesis. To assist the physician in evaluating cardiac function, we introduce an intuitive visualization of the motion data in a color-coded fashion inside a time-varying region of interest. A first clinical validation of the algorithm is performed on ultrasound data obtained from an animal study.

In Chapter 5, we investigate the feasibility of the so-called total-least-squares method to fit the local model parameters. This discussion is motivated by the observation that the least-squares estimator used in the previous chapter may be biased if the system matrix is corrupted by noise. We discuss the relation between the two estimators and also consider the so-called scaled total-least-squares method that represents a trade-off between the two different approaches.

In Chapter 6, we present an extension of the proposed motion analysis method that takes advantage of additional measurements from tissue Doppler ultrasound. We derive a corresponding bimodal motion estimation algorithm and demonstrate that the incorporation of the partial, Doppler-based velocity measurements increases the accuracy of the algorithm. We validate the approach with controlled phantom experiments and also present some clinical results that correlate well with the expert echocardiographic reading.

We then concentrate on the estimation of myocardial deformation in Chapter 7. The deformation is described in terms of strain rate, a quantity that is inherently contained in the affine motion parameters. In contrast to the state-of-the-art Doppler-based strain rate imaging, the method is not limited to the scan line directions; it rather allows to determine the two principal directions as well as the amount of deformation. We account for the physical incompressibility of the myocardium by adding a divergence-free constraint to the local velocity model. The accuracy and noise sensitivity of the approach is tested on synthetic ultrasound data. First applications to clinical echocardiograms yield reasonable qualitative results.

In Chapter 8, we validate the proposed motion analysis algorithm on a large set of B-mode echocardiograms obtained from 114 patients during routine clinical examinations. The algorithm is first compared with Doppler velocities that were available on a subset of the patients. We also conduct a detailed statistical comparison between velocities computed by the algorithm and conventional echocardiographic reading. Lastly, we investigate the diagnostic relevance of the extracted motion parameters by using a neural network-based classifier. We

also highlight the clinical value of the algorithm on exemplary case studies.

Appendix A includes additional research results that are relevant to echocardiography but have a wider range of applicability. In particular, we introduce the general concept of multiresolution moment filters. Inspired by the framework of Wavelets, we developed this multi-channel filtering scheme to compute weighted geometric local moments efficiently at dyadic scales. Beyond its application in the proposed motion analysis algorithm, we demonstrate its versatile usefulness in other areas of image analysis such as image denoising and feature extraction.

It is not by the gray of the hair that one knows the age of the heart.

— Henry Bulever

Chapter 2

Basic Heart Anatomy and Physiology

Abstract — We give a brief overview of the heart anatomy and of its normal function. We then discuss common cardiovascular diseases that lead to wall motion abnormalities. Finally, we present the conventional diagnostic assessment of myocardial dysfunction using a semi-quantitative wall motion scoring.

2.1 Introduction

The heart is a restless working muscle that provides blood circulation to the whole human body. At an average rate of 72 times per minute, the heart beats about 100 000 times per day and 30 to 42 million times per year. This corresponds to a daily pumping capacity of about 7200 liters and 2.5 million liters per year. This enormous activity is due to the heart's highly optimized structure and functioning. To provide the reader with the necessary background and to be able to understand the origin and symptoms of cardiac malfunction, we first review the heart's basic anatomy and its associated terminology in Section 2.2.1. We then discuss the normal heart physiology and electrical conduction system in Sections 2.2.2 and 2.2.3, respectively.

Several heart diseases affect the heart muscle and reduce its ability to contract normally. The resulting loss in pumping efficiency leads to an undersupply of the vital oxygen to the body, which may seriously harm other organs. The most frequent source for impaired wall motion is a process called atherosclerosis, which, in its worst case, may lead to a life threatening heart attack. This and other common diseases and their impact to cardiac function are summarized in Section 2.3.

In clinical routine, myocardial dysfunction can be diagnosed by a number of tools including physical examinations, electrocardiograms, blood tests, and, in particular, medical imaging techniques such as echocardiography. Although current imaging technologies are highly sophisticated, myocardial malfunction is usually assessed by visual inspection and subjective rating of the acquired images by experienced cardiologists. In echocardiography, the presence and extend of wall motion abnormalities is commonly classified according to the so-called wall motion score index; this conventional assessment is reviewed in Section 2.4.


2.2 Heart Anatomy and Normal Functioning

In this section, we review basic heart anatomy and physiology. We also introduce the medical terminology that will be frequently used in this thesis.


2.2.1 Heart Anatomy

The heart is a muscular organ that is located between the lungs in the middle of the chest, behind and slightly to the left of the breastbone (sternum). It is surrounded by a double-layered membrane that is called the *pericardium*. The heart has four chambers as illustrated in Fig. 2.1. The two upper chambers are called the left and right *atria*, and the lower chambers are called the left and right *ventricles*. The atria act as reservoirs for venous blood; they have a small pumping function to assist ventricular filling. The ventricles are the major pumping chambers that deliver blood to pulmonary (right ventricle) and systemic circulations (left ventricle). The heart's muscular walls are called *myocardium*. Its outer surface is called the *epicardium* and its inner lining the *endocardium*. The wall that separates the left and right atria and the left and right ventricles is called the *septum*.

Four valves ensure that the blood flows only in one direction and prevent blood from leaking backwards from one chamber to the upstream chamber

**External View:**

1. Right Coronary (RCA)
2. Left Anterior Descending (LAD)
3. Left Circumflex (LCX)
4. Superior Vena Cava
5. Inferior Vena Cava
6. Aorta
7. Pulmonary Artery
8. Pulmonary Vein

**Internal View:**

9. Right Atrium
10. Right Ventricle
11. Left Atrium
12. Left Ventricle
13. Papillary Muscles
14. Chordae Tendineae
15. Tricuspid Valve
16. Mitral Valve
17. Pulmonary Valve

Figure 2.1: Anatomy of the human heart. Anterior external view with coronary arteries and veins (top) and internal view (bottom). (Image source: <http://www.cardioconsult.com/Anatomy/>)

(valvular regurgitation). The aortic and pulmonic valves are referred to as the semilunar valves and are located at the downstream sides of the left and




Figure 2.2: Normal heart physiology during diastole and systole. The right atrium (RA) receives the de-oxygenated blood from the body that is pushed to the lungs by the right ventricle (RV). The re-oxygenated blood flows via the left atrium (LA) into the left ventricle (LV) that pumps it into the body. In addition to systolic inward motion and longitudinal shortening, the myocardium thickens considerably. The process is reversed during diastole. (Figure adapted from: <http://www.tmc.edu/thi/valves.html>)

right ventricle, respectively. The two atrioventricular valves, the mitral and tricuspid valve, are located between the atria and ventricles. The leaflets of the atrioventricular valves are connected to the papillary muscles that are, in turn, connected to the walls of the ventricles. The papillary muscles shorten during the contraction of the ventricles in order to prevent a bulging of the atrioventricular valves towards the atria that would lead to regurgitation.

2.2.2 Normal Physiology

The heart cycle is a pumping action that is divided into two alternating periods of *diastole* (relaxation) and *systole* (contraction). As indicated in Fig. 2.2, the two phases proceed as follows:

Diastole: During diastole, the atria contract and push the blood into the re-

laxed ventricles. More precisely, the de-oxygenated blood returning from the body flows from the right atrium through the tricuspid valve into the right ventricle. In parallel, the left atrium pushes the re-oxygenated blood that arrived from the pulmonary veins through the mitral valve into the left ventricle. More specifically, the diastole consists of four sub-phases: (1) isovolumetric relaxation: the ventricular pressure decreases, but its volume remains constant prior to opening of the atrioventricular valves; (2) rapid filling: after opening of the atrioventricular valves, the pressure gradient drives blood from each atrium into the respective ventricle; (3) diastasis: the ventricular filling is mainly due to direct blood flow from the veins; (4) atrial systole: the atria contract actively and complete the ventricular filling.

Systole The second part of the heart cycle, the systole, begins when the ventricles are filled and start to contract. The right ventricle pumps the oxygen-poor blood through the pulmonic valve into the pulmonary artery towards the lungs to be re-oxygenated. The left ventricle pushes the oxygen-rich blood through the opened aortic valve into the aorta and further into the body. Since the left ventricle supplies the whole body with re-oxygenated blood, it is the most dominant and strongest muscle of the heart. Its pumping capacity is achieved by a highly optimized motion pattern of the ventricular walls: as indicated in Fig. 2.2, the myocardium moves transversely towards the ventricular center and also shortens in the longitudinal direction. In addition, the ventricular walls thicken significantly to increase the ejection efficiency. A rotational “wringing” or “twisting” action around the longitudinal axis increases its efficiency even further. After the pulmonary and aortic valves have closed, the myocardium performs the reverse motion, causing the ventricles to relax. The lower pressure in the ventricles causes the tricuspid and mitral valves to open, and the cycle begins again.

2.2.3 Electrophysiology of the Heart

The coordinated contraction of the various chambers of the heart is controlled by the electrical system of the heart. The so-called *electrocardiogram* (ECG) is a recording of the electrical changes that occur in the myocardium during a cardiac cycle. A typical ECG is shown in Fig. 2.3. The electrical changes result from depolarization and repolarization of cardiac muscle fibers. The electrical signal that initiates each heart beat arises from a small structure—called the




Figure 2.3: Typical ECG during heart cycle. The P-wave corresponds to atrial depolarization followed by the QRS-complex, initializing ventricular contraction. The T-wave indicates ventricular repolarization.

sinoatrial (SA) node—that is located at the top of the right atrium; it is the heart’s natural pacemaker. The depolarization of the SA-node generates an electrical signal that rapidly spreads across the right and the left atrium and causes them to contract. This impulse is referred to as the P-wave in the ECG. The impulse is then delayed at the atrioventricular (AV) node that is located in the septal wall of the right atrium. The AV-node is connected to a bundle of special cells in the heart designed to rapidly conduct the electrical signal through the ventricles. These fibers are referred to as the bundle of His. The bundle of His branches downstream into a right and a left bundle which eventually branch out to distant ventricular tissues that are referred to as Purkinje fibers. When the cardiac impulse reaches the ventricular fibers, they rapidly depolarize giving rise to the QRS-complex in the ECG. The depolarization induces the contraction of the ventricular muscles. During contraction, the ventricular muscle fibers repolarize slowly, corresponding to the T-wave in the ECG.

2.2.4 Coronary Circulation

The heart muscle, like every other organ or tissue in the body, needs oxygen-rich blood and vital nutrients to survive. Blood is supplied to the heart by its own vascular system, called coronary circulation. As can be seen in Fig. 2.1 on page 11, the aorta branches off into two main coronary blood vessels, the so-called *coronary arteries*. These coronary arteries branch off into smaller arteries, which supply oxygen-rich blood to the entire heart muscle. The right coronary

artery (RCA) supplies blood mainly to the right side of the heart. The left coronary artery, which branches into the left anterior descending artery (LAD) and the circumflex artery (LCX), mainly supplies blood to the left side of the heart. Partial or complete occlusion of the coronary arteries is one of the most frequent cause for cardiac malfunction; it is discussed next.

2.3 Common Cardiovascular Diseases

Ischemic heart disease and myocardial infarction most commonly affect the left ventricle because of its larger size and greater need for oxygenated blood. Because this chamber supplies most of the heart's pumping power, it is very essential for normal function. Right-ventricular heart failures often occur as a result of left-sided dysfunction; therefore, most of the attention to describe and analyze myocardial malfunction is drawn to the left ventricle.

2.3.1 Coronary Artery Disease

Most frequently, wall motion abnormalities are caused by the so-called *coronary artery disease* (CAD) that is the end result of a complex process called *atherosclerosis* (hardening of the arteries). Substances such as inflammatory cells, proteins and calcium that travel in the bloodstream start sticking to the vessel walls. Fat and other substances combine to form a material called plaque that builds up and narrows the arteries (*stenosis*). When plaque narrows the inside of an artery to a point where it cannot supply enough oxygen-rich blood to meet the organ's needs, cramping of the muscle occurs; this is called *ischemia*.

The most common symptom in patients with ischemia is chest discomfort or pain, known as *angina pectoris*. A distinction is drawn between two types of angina: stable angina is caused by a temporarily reduced blood supply to the heart and usually appears as the result of exertion or intense emotions. It is relatively predictable and its intensity and frequency remains similar over long periods. The second type, unstable angina, may even occur when the patient is resting. It is severe and unpredictable and often threatens to progress to an acute myocardial infarction.

A myocardial infarction is the result of a prolonged artery occlusion lasting for more than 15 minutes. It is also referred to as a "heart attack". The areas of heart muscle that are supplied by the blocked arteries die. Since myocardial tissue does not regenerate after injury, it is replaced by scar tissue that usually inhibits contractility. The larger the affected area of the myocardium the greater

the loss of contractility will be. The detection and quantification of regional wall motion abnormalities is an important tool to evaluate the degree of ischemia and infarction.

If the occluded artery is not cleared by surgery or medical intervention, the infarction may lead to other complications, including various disturbances in the normal heart rhythm, congestive heart failure, cardiogenic shock, or thromboembolisms.

2.3.2 Other Diseases

Although coronary artery disease is the principal cause for regional wall motion disturbances, they may also originate from a variety of other cardiac dysfunctions. Among others, the most frequent diseases are:

Dilated (congestive) cardiomyopathy: The heart cavity is enlarged and stretched (cardiac dilation), but the wall-thickness remains normal. This results in a diffuse, globally restricted myocardial contraction. There may be multiple causes including alcoholic excess, various infections and sometimes heredity.

Hypertrophic cardiomyopathy: This genetic disorder causes an abnormal growing of the left ventricular wall thickness. The thickening leads to wall stiffness and might prevent the heart from relaxing during diastole.

Myocarditis: Myocarditis is an inflammatory disease of the myocardium that can result from a variety of causes such as viral infections, drugs, and hypersensitive immune reactions. It may lead to distinct regional wall motion abnormalities.

Valvular heart disease: Wall motion disturbances may also appear as late effects of mitral or aortic valve insufficiencies and aortic stenosis.

Pure septal defect: Wall motion abnormalities that are restricted to the septum appear frequently. They are usually characterized by a temporally perturbed contraction pattern, but show a normal wall thickening. In particular, the septum is called “paradoxical” if it moves outwards at early systole. This defect can be caused by overpressure in the right ventricle, left branch block or pericardial effusion and might also appear after heart surgery.

For an extensive description of heart anatomy, physiology and cardiovascular diseases we refer to [22, 23].

2.4 Conventional Diagnostic Assessment of Myocardial Malfunction¹⁷

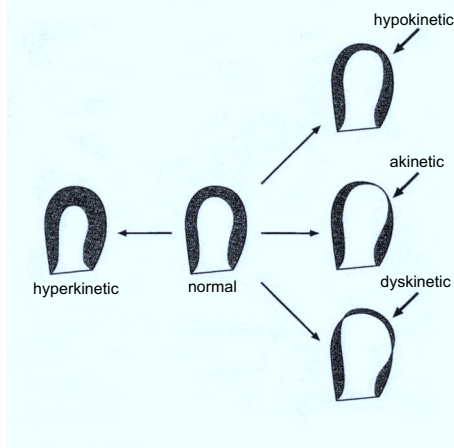


Figure 2.4: Conventional categorization of wall motion abnormalities according to their degree of severity. Shaded regions represent the course of motion of the endocardium during systole. (Figure adapted from [26])

2.4 Conventional Diagnostic Assessment of Myocardial Malfunction

The presence of regional contraction abnormalities strongly suggests the diagnosis of coronary artery disease. Several studies have shown a very good correlation between the extent of wall motion abnormalities and the size of an infarction [24].

In clinical practice, myocardial motion is usually assessed by visual inspection and manual evaluation of echocardiograms. The most common approach is based on the so-called wall motion score index. The ventricular wall is divided into regional segments, each of which is assigned a score based on the severity of the dysfunction. The segmental partitioning that is typically used in echocardiography is described in Section 3.3 later on. A widely used scoring system, developed by Kan *et al.* [25], is illustrated in Fig. 2.4. Correspondingly, the degree of infarction is ordered according to the following categories:

Hyperkinesia (-1): The non-affected, normally perfused part of the myocardium may contract excessively to compensate for the impaired contraction of ischemic or infarcted segments.

Normokinesia (0): The myocardium moves and thickens normally.

Hypokinesia (1): The affected segments move slower and thicken less than normal.

Akinesia (2): The infarcted region has completely lost its ability to contract and moves passively along with its surrounding tissue.

Dyskinesia (3): The infarcted region behaves paradoxically and bulges during systole due to the ventricular blood pressure.

Aneurysm (4): The infarcted region has become very thin, bulges extremely during systole and even remains bulged during diastole.

As clinical practice shows, this subjective, semi-quantitative assessment of myocardial dysfunction relies heavily on the observer's experience; resulting diagnoses suffer from a high intra- and interobserver variability. Thus, a more quantitative, observer-independent motion analysis is highly desirable to increase their objectivity and reproducibility.

There are similarities, but each heart is different.

— Michael DeBakey

Chapter 3

Principles and Instrumentation of Echocardiography

Abstract — In this chapter, we briefly review the role of ultrasound in cardiac imaging and introduce its physics and basic principles of image formation. We give an overview about different imaging modes and views that can be obtained from the heart by using different acquisition techniques. The reader is also familiarized with the standard technical and anatomical nomenclature that is used in clinical echocardiography.

3.1 Introduction

Various imaging modalities are available for investigating the structure and dynamics of the heart. They differ largely with respect to image quality, ease of operation, infrastructure requirements, and degree of invasiveness.

ECG-gated single photon emission computed tomography (SPECT) [27, 28] and positron emission tomography (PET) [29] are radionuclide-based methods that are mainly designed to analyze myocardial perfusion. Their ability to assess myocardial motion is limited due to their low spatial and temporal resolution.

The high infrastructure requirements such as an on-site cyclotron and the elaborate way of operation further limit its applicability in clinical routine.

X-ray computed tomography (CT) has also been applied for cardiac imaging. In particular, cine CT, also known as ultrafast CT (UFCT) or electron beam CT (EBCT), has been used for cardiac motion assessment [30, 31]. Although CT has mainly been used to screen for coronary atherosclerosis based on calcium scoring, it is becoming more and more popular for functional studies. This is due to the recent development of fast ECG-gated multislice spiral computed tomography (MSCT) [32]. This technique offers a higher in-plane resolution and thinner slice collimation but still suffers from a lower temporal resolution.

Cardiac magnetic resonance imaging (MRI) [33] has emerged as an important approach to assess myocardial anatomy. MRI offers a high contrast discrimination between soft-tissue and blood without the need of a contrast medium. MR tagging [34, 35], in particular, was introduced to analyze myocardial motion. This technique is able to generate grid lines at one point in the cardiac cycle by spatial modulation of magnetization (SPAMM). The tissue motion is then indicated by the displacement and deformation of the grid. However, it is difficult to track the tags over the complete cardiac cycle since the magnetization decays with time. Phase-encoded MRI is a different motion analysis approach that is based on measurements of MR phase changes. The phase change is used to estimate the instantaneous velocities of myocardial tissue [36]. Wide application of MR imaging remains limited by cost and the difficulty in routinely applying these MR approaches to critically-ill patients.

Echocardiography offers significant advantages over the aforementioned modalities. Reasons for its wide acceptance are its non-invasiveness, portability, ease of use, and cost effectiveness. However, echocardiography is limited by the quality of the images that can be obtained in an individual patient. The rib cage and lungs may preclude good visualization of the entire heart because sound waves cannot penetrate these structures. Three types of echocardiography are commonly used: M-mode, two-dimensional B-mode and Doppler. Each has important applications which are briefly described in this chapter.

This chapter is organized as follows: we first review the basic principles of ultrasound physics and image formation in Section 3.2. In Section 3.3, we introduce the different instrumentation techniques in clinical echocardiography; in particular, we explain the standard two-dimensional views and the corresponding cardiac sections imaged.

3.2 Basic Principles of Ultrasound

Diagnostic ultrasound employs pulsed, high frequency sound waves to image cardiac structures and its motion. The ultrasound wave is generated by an ultrasound transducer, which usually employs a piezoelectric material. An electrical pulse, generated by a pulser/receiver unit, is converted into an acoustic pressure field with some carrier frequency. The inverse piezoelectric effect is used to receive the ultrasound wave that is converted into an electrical *radio frequency* (RF) signal. The RF-signal is then processed by the ultrasound machine to create characteristic images. The transducer emits the acoustic pulses with a repetition frequency typically ranging from 0.5 to 20 kHz, based on the maximum target depth. The propagational speed of the ultrasound pulses is determined by the elasticity and density of the medium, and is nearly constant in the soft tissues of the body (approximately 1540 m/s).

The emitted ultrasound waves undergo a complex interaction with the ionified tissue; the interaction can basically be divided into the following three types:

Reflection and Refraction

Reflection or refraction of the sound wave occurs at approximately planar surfaces that have irregularities that are much smaller than the acoustic wavelength. Such surfaces are referred to as specular. Like in optics, sound waves are reflected from a specular interface at an angle that is equal to the angle of incidence. The portion of the reflected intensity depends on the difference between the sound impedances of the two adjacent tissues. The part of sound that is transmitted through the interface is refracted by an angle that depends on the ratio of the propagation velocities in the two tissues.

Excessive reflection may significantly degrade the resulting image quality. It appears mainly in cases where the epicardial and pericardial surfaces are oriented such that much of the beam energy is reflected out of the line of sight.

Scattering

The interaction of the acoustic wave with cellular structures that are much smaller than its wavelength is very different than the one with specular surfaces. According to Huygens' principle, such point targets act as scatterers that emit spherical secondary waves. Their intensity depends on the ratio of the acoustic wavelength and the scatterer radius. The superposition of all scattered waves




Figure 3.1: Block diagram of image formation of a typical ultrasound system. TGC: time-gain compensation; A/D: analog-to-digital converter; LPF: lowpass filtering.

results in a constructive and destructive interference pattern. This portion of the image is referred to as acoustic *speckle* and is the major reason for the noisy appearance of ultrasound images.

Absorption

As a sound wave passes through the tissue, it progressively loses energy. The amplitude attenuation is primarily due to the inner friction or viscosity of the tissue that transforms sound energy into other energy forms such as heat. The signal attenuation depends highly on the carrier frequency ω_c ; higher frequencies allow a better spatial resolution, but are more attenuated than lower ones and thus have less penetrating ability. Conversely, a lower frequency transducer has a greater depth of penetration but poorer resolution. The optimal carrier frequency is a trade-off between the requirements of penetration depth and image resolution. Frequencies generally used in diagnostic ultrasound range from 3.5 to 10 MHz. The spatial resolution in axial direction of a typical 5 MHz probe is about 0.3 to 0.6 mm. The lateral resolution is usually four to five times lower, i.e., 1.2 to 3.0 mm.

As illustrated in Fig. 3.1, the received electronic signals go through several steps of signal processing before they are displayed: first, the receiver amplifies the returned signal in proportion to the target depth that is given by the time required by the signal to return. This step—called time-gain compensation (TGC)—corrects the signal attenuation during the propagation of sound in the tissue. Other amplitude controls such as the user-adjustable coarse gain control may also be applied. After amplification, the radio frequency (RF) signal

is sampled by an analog-to-digital converter. Then, the carrier frequency ω_c is removed by techniques such as quadrature demodulation. In this technique, the RF-signal is multiplied by $\cos(\omega_c t)$ and $\sin(\omega_c t)$ and then lowpass filtered (LPF) to obtain the complex baseband-equivalent signal $I(t) + iQ(t)$. $I(t)$ and $Q(t)$ are the time varying inphase and quadrature components of the demodulated signal, respectively. The down-mixed signal is usually sub-sampled to reduce the amount of data. The amplitude $E(t) = \sqrt{I(t)^2 + Q(t)^2}$ is obtained by envelope detection of the complex baseband signal.

If the amplitude of the returned echo is displayed in form of gray levels, it is logarithmically compressed to adjust the dynamic range of the signal to the small grayscale range of the monitor. This non-linear mapping enhances darker gray levels at the expense of the brighter gray levels. Several techniques such as speckle suppression may then be applied to improve the quality of the displayed signal. The final amplitude corresponds to the strength of the returned echo and is presented to the physician for interpretation.

For a detailed description of ultrasound physics and its image formation, we refer to [37, 38, 39].

3.2.1 Imaging Modes

The detected echoes may be displayed in one-dimensional formats such as A-mode, B-mode, M-mode, or in the two-dimensional B-mode format. These different display modes are briefly described in the following.

Single-Beam Systems

One-dimensional ultrasound systems are designed to measure the distance from the transducer surface to reflecting structures in the heart along a pre-defined direction. The received signal can be displayed in several ways:

A-mode Amplitude mode is a one-dimensional ultrasonic display showing echo intensities along the ultrasonic beam as vertical spikes on a horizontal time axis indicating the depth of the reflectors.

B-mode This method displays the received echo as a spot whose brightness is proportional to the echo intensity. This type of display is referred to as the *brightness mode*.

M-mode The motion mode corresponds to a dynamical display of brightness mode intensities. The M-mode tracing displays distance from the transducer on the vertical axis and time on the horizontal axis. Since M-mode

employs only a single beam, it achieves a high temporal resolution which allows to analyze quick motions such as of moving heart valves.

In many cardiac applications the display of anatomic structures and spatial relationships is preferred to a one-dimensional analysis, although this is only feasible at a lower temporal resolution. The display of multiple B-mode lines in a two-dimensional image is discussed in the next section.

Two-dimensional Ultrasound Systems

In two-dimensional echocardiography, the ultrasound beams are transmitted at a series of angles to analyze a sector of a plane. At each scan line position, one ultrasound pulse is transmitted and all echoes from the surface to the deepest range are recorded before the ultrasound beam moves on to the next scan line position, where pulse transmission and echo recording are repeated. Various transducers are used in ultrasound imaging. Single element transducers must be steered mechanically over the region of interest, whereas array transducers can steer the beam electronically. In addition, array transducers are able to vary the depth of the focus plane. The dynamic beam steering and focusing is accomplished by firing the individual elements with suitable time delays. Such *phased-array* transducers are most widely utilized in cardiology since, among others reasons, their compact size accommodates the transmission between the ribs.

For display purposes, the data given in polar coordinates is transformed to the cartesian grid in a process that is known as *scan conversion*. As shown in Fig. 3.2, each pixel value on the cartesian grid is interpolated from its surrounding data in polar coordinates. This task is usually accomplished by means of nearest neighbor or linear interpolation. Recently, we have proposed a more general and comprehensive method that uses higher interpolation orders and a global smoothness constraint [40].

In the past decade, three-dimensional echocardiography has become more feasible and is emerging as a promising approach to cardiac imaging. A relatively widespread 3D-imaging technique is provided by transoesophageal echocardiography (TEE). It involves the placement of the ultrasound transducer into the oesophagus in proximity to the heart. The TEE-probe is typically mounted at the tip of a flexible endoscope and rotated at defined angle increments. Gated by respiration and the electrocardiogram (ECG), two-dimensional images are sequentially acquired. The quality of the images is superior to transthoracic




Figure 3.2: Data before scan conversion (left) and after scan conversion (right). The data is converted from cylinder coordinates (r, ϑ) to the cartesian grid (x, y) .

recordings because interferences with the chest wall and lungs are excluded. The acquired image planes may be post-processed off-line to reconstruct a conical volumetric data set [41, 42]. Since the method is semi-invasive and time consuming compared with standard transthoracic echocardiography, it is only utilized for specific indications.

Recently, development has proceeded towards transducers that consist of two-dimensional crystal matrix arrays [43]. Here, scan lines can be directed by two-dimensional phasing in any direction within a pyramidal volume and thus provide true three-dimensional imaging. However, image quality is still limited by a low spatial and temporal resolution; it is therefore not yet widely applicable in clinical practice.

3.2.2 Doppler Echocardiography

Tissue Doppler imaging (TDI) is an ultrasound modality that is based on the so-called Doppler principle [44]. TDI is a further development of conventional Doppler ultrasound that measures blood flow velocity [45]. Since myocardial tissue moves about ten times slower than blood cells, it is possible to extract regional myocardial velocities by using thresholding and filtering algorithms that

reject the echoes originating from the blood pool. TDI quantifies tissue velocities by detection of consecutive phase shifts of the ultrasound signal reflected from the myocardium. The frequency shift is due to the relative motion between the wave source and the receiver. If the reflecting tissue is moving towards the transducer, the frequency of the returning signal increases; conversely, the frequency of the returning signal decreases if the reflector moves away from the transducer.

The Doppler shift can be measured by using two different techniques: in pulsed Doppler, a single crystal is used to transmit and receive signals. The advantage of pulsed Doppler is that a specific depth of interest can be targeted. However, the maximum velocity that can be measured unambiguously is limited by half the frequency shift of the pulse-repetition frequency. If the tissue velocity exceeds this Nyquist limit, aliasing artifacts occur. Aliasing can be avoided by using continuous wave Doppler. This technique uses two crystals, one to transmit continuously and one to receive the returned signal. This allows to measure higher velocities than in pulsed Doppler. A disadvantage of this method is the lack of spatial focus, since the signal is recorded along the full length of the ultrasound beam. Consequently, velocities from different moving structures will be averaged.

Velocities can be measured along a single beam or across a whole image section. Single-beam measurements are usually displayed in the so-called spectral format, where velocities are plotted versus time in a coordinate system. If the Doppler signal is measured for an image region, it is usually color-coded and superimposed on the B-mode grayscale images in real-time. This display format is also known as *color Doppler*.

The fundamental limitation of tissue Doppler is that it only measures the axial velocity component along the scan line. In contrast to blood flow, which is essentially one-dimensional, as it is oriented along the blood vessels, the myocardium performs complex, multi-dimensional motion patterns that can be hardly assessed by Doppler. However, in Chapter 6, we will propose a novel method to integrate this partial information into a B-mode-based motion analysis algorithm to obtain more reliable, two-dimensional motion data.

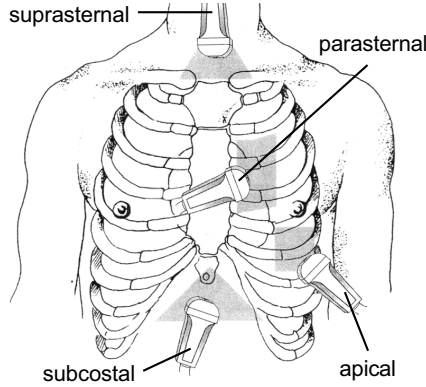


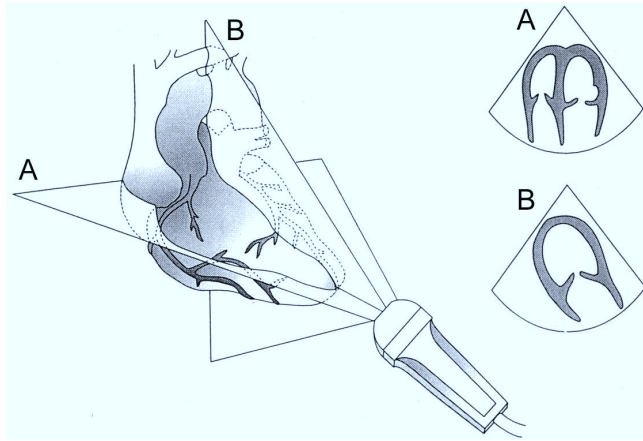
Figure 3.3: Acquisition windows for transthoracic echocardiography. (Figure adapted from [46])

3.3 Standard Views for Two-Dimensional Echocardiography


Among the different types of imaging modes mentioned above, two-dimensional echocardiography is the backbone of cardiac ultrasound. However, access to the heart is limited because the emitted sound beam cannot pass through either bone or the lungs. As illustrated in Fig. 3.3, there are a few strictly limited acoustic windows through which the heart can be interrogated transthoracically. Two of these windows, the *apical* and *parasternal* ones, permit the heart to be cut along its long or short axis, respectively. These two windows are used mainly for the analysis of ventricular function, whereas the suprasternal and subcostal windows are primarily applied for other investigations. For instance, the suprasternal window is well suited for imaging the aorta while the subcostal window allows the imaging of the interatrial septum and the inferior vena cava. In practice, the sonographer manipulates the transducer on the patient's chest to optimize image quality and position via real-time image feedback before recording each view. Not all views will have acceptable image quality, or are even obtainable, in all patients.

3.3.1 Apical Views


The two orthogonal planes that are routinely acquired from the apical window are the four- and two-chamber views. They are acquired along the heart's long axis as illustrated in Fig. 3.4a.



(a) Examining planes from apical window. A: apical four-chamber view, B: apical two-chamber view (Figure reproduced from [46])



(b) Frame of apical four-chamber view¹



(c) Frame of apical two-chamber view¹

Figure 3.4: Echocardiographic image acquisition from apical window.

¹The images are courtesy of the author's heart.

The apical four-chamber view, shown in Fig. 3.4b, is one of the most important views since it includes parts of each of the four cardiac chambers and gives a good overview of the complete heart. In particular, the section shows the interventricular and interatrial septa in the center, with the left ventricle and atrium to the right, and the right ventricle and atrium to the left.

The primary application of the two-chamber view is to assess the function of the left ventricle. By definition, only the left atrium and ventricle are imaged. This view is difficult to acquire, but it is the only one where the whole anterior (front) wall is visible. A typical example is shown in Fig. 3.4c.

3.3.2 Parasternal Views

The parasternal views, normal to the long axis of the left ventricle, are obtained by directing the ultrasound beam through a gap in the ribcage. It permits the heart to be imaged from front to back, and also in short axis.

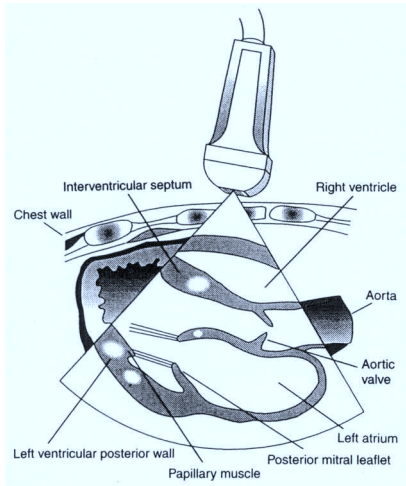
As illustrated in Fig. 3.5a and 3.5c, the parasternal long axis views are particularly useful for displaying the left atrium, the mitral valve, the inflow and outflow tracts of the left ventricle, and the aortic valve. The apex of the left ventricle is often excluded from the parasternal long axis view since it is well visualized in the apical two-chamber view.

By changing the angulation on the chest wall, it is possible to obtain any number of short axis views. The standard four views are at the level of the aortic valve (AV), mitral valve (MV), left ventricular papillary muscles (PM), and left ventricular apex (AP); their acquisition is illustrated in Fig. 3.5b. An echocardiogram obtained at the papillary muscle level is shown in Fig. 3.5d.


3.3.3 Echocardiographic Nomenclature

In this work, we follow the nomenclature of the American Society of Echocardiography [47] to denote the wall segments that are imaged in the different cross-sectional views. As shown in Fig. 3.6, this nomenclature divides the left ventricle vertically into three parts: the apical (top), the mid-cavity and the basal (bottom) segment.


Circumferentially, the basal and mid segments are further divided into anterior (front wall), lateral, posterior, inferior (back wall), septal, and anteroseptal parts, respectively. The apical segment is only divided into four circumferential parts (anterior, septal, lateral, inferior) because its size is much smaller than the one of the mid and basal segments. All 16 segments are shown in Fig. 3.7 in the form of a “bull’s eye”-plot.




(a) Examining plane in parasternal long axis view¹



(b) Examining planes in parasternal short axis views: A: aortic valve level, B: mitral valve level, C: papillary muscle level, D: apical level¹



(c) Frame of parasternal long axis view²



(d) Frame of parasternal short axis view at papillary muscle level (C)²

Figure 3.5: Echocardiographic image acquisition from parasternal window.

¹Figure reproduced from [46].

²The images are courtesy of the author's heart.




Figure 3.6: 16-segment model of the left ventricular myocardium according to the American Society of Echocardiography. Shaded regions indicate the blood supply by the corresponding coronary arteries, respectively.

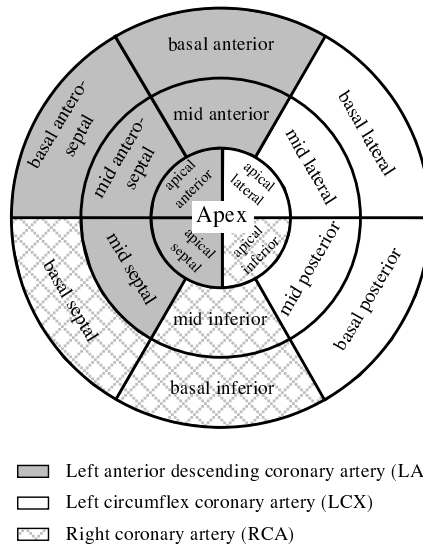


Figure 3.7: “Bulls eye”-plot of the 16-segment model of the left ventricular myocardium according to the American Society of Echocardiography. Shaded regions indicate the blood supply by the corresponding coronary arteries, respectively.

3.4 Conclusions

The basic principles of ultrasound and its instrumentation have been summarized. Echocardiography yields information on the heart's anatomy and dynamics non-invasively that can be displayed in various fashions. Partial, quantitative motion information can be obtained by using the Doppler principle. The greatest limitation of echocardiography is the lack of an adequate access to the heart. This means that a fair amount of technical skill is required to obtain images of acceptable quality.

With this review, we hope that we have helped the reader to develop a better appreciation of the problems and challenges that one is faced with when developing an automated and quantitative analysis of echocardiograms.

It is not the size of a man but the size of his heart that matters.

— Evander Holyfield

Chapter 4

Motion Analysis from B-mode Echocardiograms

Abstract — We present a new optical-flow-based method for estimating heart motion from two-dimensional B-mode echocardiograms. To account for typical heart motions such as contraction/expansion and shear, we analyze the images locally by using a local affine model for the velocity in space and a linear model in time. The regional motion parameters are estimated in the least-squares sense inside a sliding spatio-temporal B-spline window. Robustness and spatial adaptability is achieved by estimating the model parameters at multiple scales within a coarse-to-fine multiresolution framework. We use a Wavelet-like algorithm for computing B-spline-weighted inner products and moments at dyadic scales to increase computational efficiency. To characterize myocardial inward/outward motion and to simplify the detection of myocardial dysfunction, the radial component of the velocity with respect to a reference point is color-coded and visualized inside a time-varying region of interest (ROI). The algorithm is first validated on synthetic data sets that simulate a beating heart with a speckle-like appearance of echocardiograms. The ability to estimate motion from real ultrasound sequences is demonstrated by a rotating phantom experiment. The method is also tested on a set of *in vivo* echocardiograms from an animal study.

This chapter is largely based on our paper [48].

4.1 Introduction

Echocardiography is an effective imaging modality that enables clinicians to study the shape, size and dynamics of the heart. It is non-invasive, relatively inexpensive and fairly simple to use. The analysis of ventricular motion, in particular, provides an efficient mean to evaluate the degree of ischemia and infarction [49, 37, 50]. In clinical practice, the analysis mainly relies on visual inspection or manual measurements by experienced cardiologists. Manual methods are tedious and time-consuming and visual assessment leads to qualitative and subjective diagnoses that suffer from a considerable inter- and intra-observer variability. Therefore, an automated, computer-based analysis is highly desirable to obtain more objective and quantitative diagnoses.

Several approaches have been proposed to quantify heart motion from 2D echocardiograms. One approach consists in segmenting and tracking myocardial borders using active contours [1, 2, 9] or active appearance models [11]. However, motion information is only obtained for myocardial borders which are often poorly defined, especially in the case of the epicardial border. A different approach is to estimate motion for the entire image content; optical-flow methods, in particular, have led to promising results: Mailloux *et al.* [17] extended the optical-flow algorithm of Horn and Schunk [16] by adding a linearity constraint to the motion field. Zini *et al.* [18] added an additional incompressibility constraint. A popular optical-flow algorithm is the Lucas-Kanade method [19], which estimates the motion locally, assuming that the velocity field is constant within a window. This method was applied to ultrasound by Chunke *et al.* [20]. An evaluation of different optical-flow methods applied to echocardiograms can be found in [21].

We propose a novel optical-flow-based algorithm that is tuned for the analysis of ventricular wall motion from dynamic B-mode echocardiograms. Inspired by the Lucas-Kanade method [19], we use a local motion model inside a sliding spatio-temporal window. We use a local affine model for the velocity in space that allows to describe typical heart motions such as rotation, contraction/expansion and shear. The motion parameters are estimated in the weighted least-squares sense inside the sliding spatio-temporal window. Instead of working at a fixed scale, we consider estimation windows of dyadic sizes and develop a multiresolution strategy to improve the estimation of large motions and to reduce the sensitivity to noise. We introduce a B-spline weighting scheme that has important computational advantages for multi-scale processing. In particular, we develop a wavelet-like multiresolution implementation. We also add a temporal linear component to our motion model. This leads to a more robust

motion estimation, combining information from multiple frames. It also yields an additional acceleration parameter—a useful indicator of heart dynamics.

Since the estimated velocity data itself is not of direct use for the clinician, it is processed and visualized properly to facilitate the diagnosis. In order to focus on the relevant regions of the heart, the motion information is only displayed inside a user-defined region of interest (ROI) that typically corresponds to the myocardium. In order to follow the movement of the myocardium, the ROI-contour is automatically tracked in time by using the estimated velocity field. Robustness of the tracking process is achieved by using a spline representation of the ROI-contours that is fitted in the least-squares sense to the estimated motion field. To extract myocardial contractility information from the overall motion field, we compute the radial component of the estimated velocity field with respect to the ROI-centroid (center of gravity). This radial velocity component is then color-coded and superimposed on the original image sequence in a semi-transparent fashion that is similar to tissue Doppler imaging. The color display allows a more intuitive and simplified identification of regions with abnormal motion patterns. As additional information, the motion trajectory of the ROI-centroid, the principal axes of inertia and the area-size of the ROI can also be displayed.

This chapter is organized as follows. We describe the motion estimation algorithm in Section 4.2. The ROI-tracking and radial motion visualization is presented in Section 4.3. The algorithm is validated in Section 4.4: first, we apply it to synthetic data that simulates the characteristics of a beating heart, and compare it with alternative motion estimation approaches. In particular, we test its robustness by simulating the ultrasonic image formation to generate characteristic speckle noise. Second, we demonstrate the ability of the proposed method to analyze motion from real echocardiograms by performing an experiment with a tissue mimicking phantom that is described in Section 4.4.2. Finally, in Section 4.4.3, we validate the algorithm on a set of clinical echocardiograms obtained from an animal study. Data sets that were acquired before and after an artificially induced infarction were analyzed and compared.

4.2 Local-Affine, Multi-Scale Motion Estimation

Several methods have been proposed to quantify heart motion from two-dimensional echocardiograms. Special attention has been paid to optical-flow

methods which have led to promising results [21]. Gradient-based optical-flow estimation relies on the assumption that the intensity of a particular point in a moving pattern does not change with time. Let $I(x, y, t)$ denote the intensity of pixels at location $\mathbf{r} = (x, y)$ and time t in an image sequence. The constant intensity assumption can be expressed as [16]

$$I_x(\mathbf{r}, t) u(\mathbf{r}, t) + I_y(\mathbf{r}, t) v(\mathbf{r}, t) + I_t(\mathbf{r}, t) = 0, \quad (4.1)$$

where I_x , I_y and I_t denote the spatial and temporal derivatives of the image intensity. The velocities u and v are, respectively, the x - and y -components of the optical-flow that we wish to estimate. Since (4.1) is a single equation in two unknowns u and v , it cannot be solved uniquely without introducing additional constraints.

4.2.1 Local Affine Velocity in Space-Time

A popular optical-flow algorithm is the Lucas-Kanade method [19], which was applied to ultrasound in at least two studies [21, 20]. The method estimates the motion locally assuming it to be constant within a spatial window. Since typical heart motions are given by rotation, expansion, contraction, and shear, we use a local affine model for the motion in space. Additionally, we introduce a linear model for the velocity along the time direction. This allows us to capture local accelerations in time better than by a locally constant model. Another advantage is that we can base our estimation on multiple frames around a given time point which is more robust than using only two frames as many classical optical-flow methods do. Let (x_0, y_0, t_0) denote the center of a small spatio-temporal image region Ω . Then the spatio-temporal-affine model is defined as

$$\begin{pmatrix} u(x, y, t) \\ v(x, y, t) \end{pmatrix} = \begin{pmatrix} u_0 \\ v_0 \end{pmatrix} + \begin{pmatrix} u_x & u_y & u_t \\ v_x & v_y & v_t \end{pmatrix} \begin{pmatrix} x - x_0 \\ y - y_0 \\ t - t_0 \end{pmatrix}. \quad (4.2)$$

The parameters u_0 and v_0 correspond to the velocity at the center point (x_0, y_0, t_0) ; u_x , u_y , u_t , v_x , v_y , and v_t are the first order spatial and temporal derivatives of u and v , respectively. The derivatives are assumed to be constant within the local neighborhood.

An attractive feature of the local affine model is that it gives also access to local strain rate information. The sub-matrix of spatial derivatives

$$\mathbf{J} = \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix}$$

describes the local variation of the velocity field in space which corresponds to myocardial tissue deformation. The computation of local strain rate parameters and its visualization is discussed in detail in Chapter 7.

We estimate the motion components at (x_0, y_0, t_0) by minimizing the weighted least-squares criterion

$$\int_{\mathbb{R}^3} w(x - x_0, y - y_0, t - t_0) \left(I_x u + I_y v + I_t \right)^2 dx dy dt. \quad (4.3)$$

The symmetric window function w gives more weight to constraints at the center of the local spatio-temporal region than to those at the periphery. A well suited window function is $w(x, y, t) = \beta^n(x)\beta^n(y)\beta^n(t)$, where β^n is the symmetrical B-spline of degree $n \in \mathbb{N}$ [51]. B-splines rapidly converge to Gaussians when their degree increases which ensures isotropy of the window in multiple dimensions. Varying the B-spline degree also allows to change the size of the window function. Additionally, the B-splines satisfy a two-scale equation which leads to an efficient computation of B-spline-weighted inner products at dyadic scales by using a wavelet-like algorithm (cf. Section 4.2.2).

By differentiating (4.3) with respect to each of the eight unknown parameters, we obtain the so-called Normal Equations $\mathbf{A}^T \mathbf{A} \mathbf{v} = \mathbf{A}^T \mathbf{b}$ in terms of local moments of the spatial and temporal derivatives of I , as defined in (4.5) on page 40. These symmetric linear systems are small (8×8) and can be solved efficiently. The system-coefficients at a given position (x_0, y_0, t_0) are of the form

$$m_{p,q,r}(x_0, y_0, t_0) = \int_{\mathbb{R}^3} w(x - x_0, y - y_0, t - t_0) \times \quad (4.4) \\ (x - x_0)^p (y - y_0)^q (t - t_0)^r f(x, y, t) dx dy dt,$$

where $0 \leq p + q + r \leq 2$ and f represents the functions $I_x^2, I_y^2, I_x I_y, I_x I_t$, or $I_y I_t$. Note that we do not compute the system components \mathbf{A} and \mathbf{b} explicitly, which would contain row-wise the weighted optical-flow constraint for each pixel in the window.

4.2.2 Choice of Window Size—Multiresolution Moments Computation

The optimal size of the window function w depends on the underlying motion field and is thus not known *a priori*. In regions where the displacement field is essentially homogenous in space, the motion can be well fitted by the affine

$$\mathbf{A}^T \mathbf{A} = \begin{pmatrix} \langle w, I_x^2 \rangle & \langle w, I_x I_y \rangle & \langle w, x I_x^2 \rangle & \langle w, y I_x^2 \rangle & \langle w, t I_x^2 \rangle & \langle w, x I_x I_y \rangle & \langle w, y I_x I_y \rangle & \langle w, t I_x I_y \rangle \\ \cdot & \langle w, I_y^2 \rangle & \langle w, x I_x I_y \rangle & \langle w, y I_x I_y \rangle & \langle w, t I_x I_y \rangle & \langle w, x I_y^2 \rangle & \langle w, y I_y^2 \rangle & \langle w, t I_y^2 \rangle \\ \cdot & \cdot & \langle w, x^2 I_x^2 \rangle & \langle w, xy I_x^2 \rangle & \langle w, xt I_x^2 \rangle & \langle w, x^2 I_x I_y \rangle & \langle w, xy I_x I_y \rangle & \langle w, xt I_x I_y \rangle \\ \cdot & \cdot & \cdot & \langle w, y^2 I_x^2 \rangle & \langle w, yt I_x^2 \rangle & \langle w, xy I_x I_y \rangle & \langle w, y^2 I_x I_y \rangle & \langle w, yt I_x I_y \rangle \\ \cdot & \cdot & \cdot & \cdot & \langle w, t^2 I_x^2 \rangle & \langle w, xt I_x I_y \rangle & \langle w, yt I_x I_y \rangle & \langle w, t^2 I_x I_y \rangle \\ \cdot & \cdot & \cdot & \cdot & \cdot & \langle w, x^2 I_y^2 \rangle & \langle w, xy I_y^2 \rangle & \langle w, xt I_y^2 \rangle \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \langle w, y^2 I_y^2 \rangle & \langle w, yt I_y^2 \rangle \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \langle w, t^2 I_y^2 \rangle \end{pmatrix} \quad (4.5)$$

$$\mathbf{v} = (u_0, v_0, u_x, u_y, u_t, v_x, v_y, v_t)^T$$

$$\mathbf{A}^T \mathbf{b} = - \left(\langle w, I_x I_t \rangle, \langle w, I_y I_t \rangle, \langle w, x I_x I_t \rangle, \langle w, y I_x I_t \rangle, \langle w, t I_x I_t \rangle, \langle w, x I_y I_t \rangle, \langle w, y I_y I_t \rangle, \langle w, t I_y I_t \rangle \right)^T$$

model within a large observation window. On the other hand, smaller windows will be better suited for processing areas where the motion pattern varies rapidly, which is for example the case at myocardial boundaries. Our strategy is thus to use a window function that is as large as possible such that the local motion model is not violated. Estimating model parameters inside larger windows increases the robustness to noise because the parameters are averaged over a larger sample size.

To find the most promising solution at a given position (x_0, y_0, t_0) , we compute first of all the local model parameters within windows of different sizes. More specifically, we use window functions that are dilated in space by dyadic factors $2^j, j \in \mathbb{N}$; i.e.,

$$w^{(j)}(x - x_0, y - y_0, t - t_0) = w\left(\frac{x - x_0}{2^j}, \frac{y - y_0}{2^j}, t - t_0\right). \quad (4.6)$$

The coefficients of the local Normal Equations (4.5) are computed for different window sizes and the systems are then solved for the model parameters. A solution is only admissible if the condition number of the corresponding linear system is below a predefined limit. Its size corresponds to the level above which a linear system is considered to be singular. In the experiments, we used values in the range of 2000 to 5000. A motion estimate is set to zero if the local average of the time derivative $\langle w^{(j)}, I_t^2 \rangle$ is smaller than a predefined noise level. For each pixel, we chose the final solution from that scale for which the normalized residual error

$$e = \frac{\|\mathbf{A}\mathbf{v} - \mathbf{b}\|_{l_2}^2}{\|w^{(j)}\|_{l_1}} \quad (4.7)$$

is minimal. Since $\|\mathbf{A}\mathbf{v} - \mathbf{b}\|_{l_2}^2 = \mathbf{b}^T \mathbf{b} - \mathbf{v}^T \mathbf{A}^T \mathbf{b}$, the residual of the least-squares solution can be computed efficiently by using the right hand side of the Normal Equations $\mathbf{A}^T \mathbf{A} \mathbf{v} = \mathbf{A}^T \mathbf{b}$ and by additionally computing the inner product $\mathbf{b}^T \mathbf{b} = \langle w^{(j)}, I_t^2 \rangle$. The matrix \mathbf{A} needs not to be known explicitly.

To decrease the computational burden, we alternatively place the dilated window functions $w^{(j)}$ on a sub-sampled grid spaced at distances of 2^j in each spatial dimension. Since in this constellation, the B-spline windows satisfy the partition of unity—i.e., they sum up to a constant—every pixel is equally covered. Motion vectors at intermediate grid points are obtained by using linear interpolation.

Computing inner products with large windows is computationally expensive; however—inspired by the framework of wavelets [52, 53]—we developed a

Table 4.1: Two-scale filters $h_{p,k}$ up to order $p = 2$ for β^3 .

l	-2	-1	0	1	2
$h_{0,0}(l)$	$\frac{1}{8}$	$\frac{4}{8}$	$\frac{6}{8}$	$\frac{4}{8}$	$\frac{1}{8}$
$h_{1,0}(l)$	$-\frac{1}{4}$	$-\frac{2}{4}$	0	$\frac{2}{4}$	$\frac{1}{4}$
$h_{1,1}(l)$	$\frac{1}{8}$	$\frac{4}{8}$	$\frac{6}{8}$	$\frac{4}{8}$	$\frac{1}{8}$
$h_{2,0}(l)$	$\frac{1}{2}$	$\frac{1}{2}$	0	$\frac{1}{2}$	$\frac{1}{2}$
$h_{2,1}(l)$	$-\frac{1}{2}$	$-\frac{2}{2}$	0	$\frac{2}{2}$	$\frac{1}{2}$
$h_{2,2}(l)$	$\frac{1}{8}$	$\frac{4}{8}$	$\frac{6}{8}$	$\frac{4}{8}$	$\frac{1}{8}$

multi-channel, Mallat-like algorithm to compute the matrix coefficients (4.4) recursively for successively larger window sizes. This multiresolution framework is applicable in both the sub-sampled and the non-sub-sampled “à trous” regime. In the sub-sampled case, the local moments at successive scales can be computed iteratively as:

$$\begin{aligned}
 & m_{p,q,r}^{(j+1)}(x_0, y_0, t_0) \\
 &= \sum_{k=0}^p \sum_{l=0}^q \left(h_{p,k}^{(j)} h_{q,l}^{(j)} * m_{k,l,r}^{(j)} \right) (2x_0, 2y_0, t_0) \\
 &= \sum_{k=0}^p \sum_{l=0}^q \sum_{m,n} h_{p,k}^{(j)}(m) h_{q,l}^{(j)}(n) m_{k,l,r}^{(j)}(2x_0 - m, 2y_0 - n, t_0).
 \end{aligned} \tag{4.8}$$

The one-dimensional two-scale filters

$$h_{p,k}^{(j)} = 2^{j(p-k)} h_{p,k} \tag{4.9}$$

are scaled versions of basic filters $h_{p,k}$ and are applied separately in the x - and y -directions. The corresponding two-scale filters $h_{p,k}$ for the cubic B-spline ($n = 3$) up to order $p = 2$ are given in Table 4.1. Practically, this means that the moments are first computed within the smallest window (initialization) and that their counterparts at successive coarser scales are obtained by multi-channel filtering and summation as in (4.8). For a detailed discussion of multiresolution moment filters, including proofs and further applications, we refer to Appendix A.

4.2.3 Coarse-To-Fine Multi-Scale Strategy

Estimating large motions of moving patterns that contain high frequencies may lead to aliasing artifacts. Due to the sampling theorem [54], the sampling frequency of the moving signal has to be at least twice as large as the spatial cut-off frequency ν_0 (Nyquist limit). In other words, the displacement $u\Delta t$ is limited by one-half of the minimum occurring wavelength $\lambda_0 = 1/\nu_0$; i.e.,

$$u\Delta t \leq \frac{\lambda_0}{2},$$

where u denotes the velocity of the pattern and Δt is the temporal sampling interval. If we assume that the signal is band-limited with the maximum spatial frequency $\nu_0 = 1/2$ and if we normalize the temporal sampling interval to $\Delta t = 1$, then the maximum displacement that can be estimated unambiguously is one pixel per frame.

To be able to estimate larger motions, we apply a coarse-to-fine strategy in space. As sketched in Fig. 4.1, we compute an image pyramid for each frame in the image sequence. In particular, we use a spline-based least-squares pyramid [55] with dyadic scale progression. At each pyramid level $0 \leq k \leq K$, the original image $I^{(0)}(x, y) = I(x, y)$, at level $k = 0$, is approximated by the spline model

$$I^{(k)}(x, y) = \sum_{m,n} c^{(k)}(m, n) \beta^n \left(\frac{x}{2^k} - m, \frac{y}{2^k} - n \right). \quad (4.10)$$

At each level, the number of spline coefficients $c^{(k)}$ is reduced by the factor 2 in each dimension, resulting in a successively coarser image approximation. The spline coefficients at level k are determined such that the approximation error

$$\left\| I - I^{(k)} \right\|_{L_2}$$

becomes minimal. The required partial derivatives of $I^{(k)}$ at each level k can be computed analytically from the B-spline representations (4.10) [51].

Our coarse-to-fine multiresolution strategy for motion estimation works as follows. Starting at the coarsest pyramid level $k = K$, the motion vectors at each grid position are determined by locally choosing the most appropriate window size as described in Section 4.2.2. The motion vectors are then transferred to the next finer resolution level $k - 1$ as initial estimates. The velocities are upscaled to the resolution level $k = 0$ by multiplying the vectors with 2^k . Values on intermediate grid positions are computed by using linear interpolation. If




Figure 4.1: Coarse-to-fine multi-scale strategy in space. N_x and N_y are the x - and y -dimensions of the images at level $k = 0$, respectively. Larger motions are estimated at coarser scales while smaller motions are refined at finer scales. At each grid position the most appropriate window size is chosen based on a *a posteriori* figure of merit.

the upscaled coarser-scale motion vector does not exceed the level-dependent length 2^{k-1} (*a priori* test), it is re-estimated in the same fashion as described above. An initial, coarser-scale estimate is replaced only, if the length of the re-estimated vector does not exceed the level-dependent limit 2^{k-1} (*a posteriori* test). Having completed the finest pyramid level, we fit a spatio-temporal B-spline model to the discrete output to obtain a global, continuous representation of the velocity field.

Note that the present strategy is a refinement of our initial approach published in [48]. The experimental results presented in this chapter have been computed using a single pyramid level ($k = 0$). The sequences were pre-smoothed instead to reduce the aliasing problem. The experiments in the subsequent chapters were performed using the refined multi-scale strategy as described above.

4.3 ROI-Tracking and Motion Visualization

Since the motion is estimated for the complete sequence, the information of interest needs to be extracted and displayed in a proper way to facilitate the identification of wall motion abnormalities. Therefore, we compute radial velocity information from the estimated velocity field and visualize it inside a time-varying region of interest (ROI). For each frame, the ROI is defined by a set of closed spline curves (typically, the inside and outside of the myocardium). For a given time t , each spline curve is parameterized as $\mathbf{r}_t(\tau) = (x_t(\tau), y_t(\tau))$ in terms of the variable τ . In particular, we represent $\mathbf{r}_t(\tau)$ as a linear combination of B-spline basis functions [51]:

$$\mathbf{r}_t(\tau) = \sum_{l=-\infty}^{\infty} \mathbf{c}_t(l) \beta^n \left(\frac{\tau}{m} - l \right). \quad (4.11)$$

The integer n denotes the degree of the B-spline and $\mathbf{c}_t(l)$ denotes a sequence of vector spline coefficients given by $\mathbf{c}_t(l) = (c_t(l), d_t(l))$. The basis functions are dilated and shifted by some integer m meaning that spline knots on the curve are located at distance m . Since the curves are closed, the sequence of spline coefficients \mathbf{c}_t is periodic with some period length K . To track a curve in time, we compute a series of sample points $\mathbf{r}_t(k)$, $k = 0, \dots, mK - 1$, at integer distance on the curve. For these sample points, we compute the displacement from the continuous spline representation of the velocity field. By adding the displacements to the current position of the sample points, we obtain their position

$$\hat{\mathbf{r}}_{t+1}(k) = \mathbf{r}_t(k) + (u(\mathbf{r}_t(k), t), v(\mathbf{r}_t(k), t)) \quad (4.12)$$




Figure 4.2: ROI-tracking by least-squares B-spline fitting to the estimated motion field.

in the next frame $t + 1$ as shown in Fig. 4.2. To obtain a robust tracking, we approximate the mK displaced sample points by the spline model (4.11) that has K degrees of freedom, i.e., the spline coefficients $\mathbf{c}_{t+1}(l)$, $l = 0, \dots, K-1$. These spline coefficients are determined by minimizing the least-squares approximation error

$$\varepsilon^2 = \sum_{k=0}^{mK-1} \|\mathbf{r}_{t+1}(k) - \hat{\mathbf{r}}_{t+1}(k)\|_{l_2}^2. \quad (4.13)$$

The choice of the parameter m determines the trade-off between closeness to the sample points and smoothness of the spline curve. An efficient filter-based approach to compute the spline coefficients is described in [56, 57]. The set of obtained spline curves then defines the ROI at time $t + 1$. To initialize the tracking process, the observer outlines the ROI by placing landmarks on the first frame of the image sequence. The user-defined landmarks are then used to obtain the initial spline curve. The corresponding spline coefficients $\mathbf{c}_0(l)$ in (4.11) are determined such that the curve exactly interpolates the landmarks ($m = 1$).

Having computed the ROI for each frame, we compute the centroids (centers of gravity), the areas, and the principal axes of inertia of the ROIs. These parameters can be computed exactly from the spline representation of the contours

[58]. The velocity field of each frame is projected onto the radial direction defined with respect to the centroid of the ROI of a particular frame (end-diastolic or end-systolic). Optionally, the velocity can also be projected onto the floating centroid of the ROI. In this case, the global translational motion of the ventricle is compensated by subtracting the velocity of the floating centroid from the estimated overall velocity. The choice of the optimal centroid (fixed vs. floating) depends on the conditions under which the heart is imaged. Under conditions where excessive translation of the heart occurs, such as might be the case after cardiac surgery, a floating centroid approach would be more appropriate.

Inside the time-varying ROI, we color-code the radial velocity using a similar color map as in tissue Doppler imaging (TDI). The color-coded velocity is then superimposed on the underlying image sequence in a semi-transparent fashion. The color display corresponds directly to myocardial inward/outward motion and allows to identify regions of abnormal motion more easily. The motion trajectory of the centroids, the principal axes of inertia and area-size of the ROIs can also be displayed. The movement of the ROI-centroid allows to identify and quantify an underlying rigid translation of the myocardium. The motion of the principal axes of inertia allows to identify a global ventricular rotation (heart-twisting).

To improve the contour tracking accuracy further, the proposed method may be combined with active contour techniques. However, most of these techniques are based on the detection of strong edges or ridges which do not necessarily correspond to the expert-defined ROI. Furthermore, the transition between the epicardial border and the surrounding tissue is usually very smooth and requires a more adapted definition of cardiac borders. The combination with trainable shape models may also improve the performance, but requires a large set of normalized training examples that also includes a variety of cases with abnormal motion. We did not consider these options here.

4.4 Numerical Results

For validation purposes, the algorithm was tested in three different ways. First, it was applied to synthetic data and compared to other motion estimation methods. Second, we analyzed echocardiograms of a rotating phantom. Third, the algorithm was applied to a set of clinical echocardiograms.

4.4.1 Application to Synthetic Data

A quantitative analysis of the performance of the algorithm is done on synthetic sequences for which the exact motion field is known.

Modeling of Echographic Texture

In order to simulate the appearance of clinical echo images, we use the linear image formation model of [59] and [60]. The model assumes that the echographic radio-frequency signal $RF(x, y)$ can be described by a spatial convolution between the system point spread function (PSF), $h(x, y)$, and the impulse response of the tissue $N(x, y)$:

$$RF(x, y) = \int_{\mathbb{R}^2} N(r, s)h(x - r, y - s) dr ds. \quad (4.14)$$

The tissue $N(x, y)$ is modeled as a collection of point-scatterers with 2D normally distributed random echogenicity. The speed of sound throughout the tissue is assumed to be constant. The PSF is assumed to be space invariant with zero attenuation in the case of an adequately adjusted gain control. In particular, we use the Gabor function

$$h(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} e^{i\omega_0 x} e^{-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2}\right)}, \quad (4.15)$$

where σ_x and σ_y correspond to the axial and lateral dimension of the PSF, respectively. The parameter ω_0 denotes the acoustic spatial frequency in axial direction. In order to simulate the radial propagation of the ultrasound beam, we divide the image plane into K radial segments along the directions $\vartheta_k = k\pi/K$, $k = 0, \dots, K - 1$. The oriented filter $h(x', y')$ in each segment is obtained from the basic filter via the coordinate transformation

$$x' = \cos(\vartheta_k)x + \sin(\vartheta_k)y \quad (4.16)$$

$$y' = -\sin(\vartheta_k)x + \cos(\vartheta_k)y. \quad (4.17)$$


The resulting B-mode images are given by the modulus (envelope) of the complex echographic radio-frequency signal. The B-mode image gray level histogram has a Rayleigh distribution which is also known as “fully formed” speckle. In this experiment, we have used the parameter set $\sigma_x = 2.0$ pixels, $\sigma_y = 2.8$ pixels, $\omega_0 = \pi/2$, and $K = 180$.

Modeling of Heart Motion


In this study, we simulate an apical short axis view of a heart. For simplicity, we consider a circular heart model and simulated a periodic displacement field that maintains the area of the myocardium constant. This is achieved by applying a radial displacement field with a magnitude decreasing with the distance from the center. This displacement field is cosine modulated in time to simulate expansion and contraction. This kind of motion reflects the wall-thinning and thickening of a beating heart during diastole and systole. Additionally to the myocardial excursion, the heart model is translated towards the upper right direction to simulate an underlying rigid motion. The known motion field is applied to the initial point scatterer image to obtain a sequence of warped images. The dynamic model was set such that the maximum velocity per frame is 2.0 pixels. In order to simulate additional changes of the scatterer echogenicity that are not caused by motion, we added normally distributed zero-mean white noise to each frame of the warped image sequence. This models, for example, echogenicity changes that are caused by blood flow. This is a simplified, not entirely realistic noise model, but allows us to test the robustness of the algorithms. The final echo sequence is computed by applying the image formation process described above to each frame of the perturbed point scatterer images. As an example, one frame of a 32-frame test sequence is shown in Fig. 4.3a. In this case, the signal-to-noise ratio of the point-scatterer images is 13.98 dB.

Motion Estimation Results

The applied motion field was estimated by the algorithm from the simulated B-mode images. As an example, the corresponding estimated velocity field of Fig. 4.3a is shown in Fig. 4.3b. The performance of the algorithm was tested on sequences that were obtained from point scatterer images of different signal-to-noise ratios. The method was compared with the Lucas-Kanade [19] and the Horn-Schunk [16] optical-flow methods. For the Lucas-Kanade method, we have used the same parameter set as proposed in [61]. For the Horn and Schunk method, we performed at most 250 steps of the iterative algorithm and used a range of regularization parameters $1 \leq \alpha \leq 150$. All sequences were prefiltered in space with a Binomial filter of variance $\sigma_s^2 = 1.0$. The experiments were performed on a 400-MHz Macintosh G4 computer. The computation time of the proposed algorithm applied to the test sequences (32 (256 × 256)-frames) was 44 seconds. To assess the performance of these algorithms, we used two error measures. The first is the angular error between the estimated velocity $\hat{\mathbf{v}}$



(a) Frame of synthetic sequence



(b) Estimated velocity field during diastole

Figure 4.3: Synthetic ultrasound sequence and corresponding estimated velocity field.

Table 4.2: Error measures of different optical-flow methods

Method	Average angular error $\bar{\theta}$ [°]	Average amplitude error $\bar{\varepsilon}$ [%]
Lucas-Kanade	9.55 ± 16.93	17.4 ± 73.8
Horn-Schunk		
$\alpha = 15$	7.77 ± 13.86	20.3 ± 59.8
$\alpha = 75$	3.90 ± 7.03	16.81 ± 21.6
$\alpha = 135$	4.05 ± 5.92	29.6 ± 18.1
Spatio-Temporal-Affine	2.43 ± 6.10	5.9 ± 15.2

and the exact velocity \mathbf{v} which is given by

$$\theta = \arccos \frac{\langle \mathbf{v}, \hat{\mathbf{v}} \rangle}{\|\mathbf{v}\|_{l_2} \|\hat{\mathbf{v}}\|_{l_2}}. \quad (4.18)$$

This error measure does not depend on the magnitude of the motion vectors, but describes the directional error of the motion vectors. The second is the relative error of the velocity magnitude which is calculated as

$$\varepsilon = \frac{|\|\mathbf{v}\|_{l_2} - \|\hat{\mathbf{v}}\|_{l_2}|}{\|\mathbf{v}\|_{l_2}}. \quad (4.19)$$

The mean errors $\bar{\theta}$ and $\bar{\varepsilon}$ and their corresponding standard deviations were computed by averaging over all pixels in the image sequence.

To test the accuracy of the different methods, they were first applied to echo sequences that were obtained from point scatterer images without additional noise. The resulting errors are summarized in Table 4.2. The worst method is the Lucas-Kanade method. This is due to the fact that the diverging character of the motion field is not compatible with the local constancy-assumption and that it uses only a single window size. The method of Horn and Schunk performs better with respect to the angular error measure. Table 4.2 shows the lowest possible angular error for the Horn and Schunk method that was obtained for the regularization parameter $\alpha = 75$. However, the regularization parameter depends on the image sequence and the best choice is not known *a priori*. Moreover, an increase of the regularization parameter leads to biased velocity

estimates towards lower values which is reflected by the relatively high amplitude error $\bar{\varepsilon} = 29.6\% \pm 18.1\%$ for $\alpha = 135$. Low values for α reduce the bias, but lead to an increased directional error. The spatio-temporal-affine method performs best since it is well adapted to the underlying motion pattern.

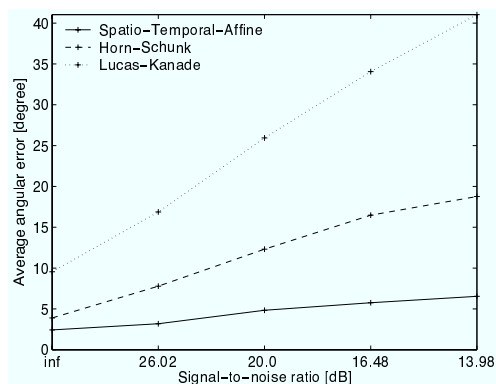
The robustness of the algorithms in the presence of additional scatterer noise is illustrated in Fig. 4.4. The algorithms were applied to echo sequences that were obtained from point scatterer images of signal-to-noise ratios varying from 26.02 to 13.98 dB. Fig. 4.4a shows the average angular errors of the three methods. For each noise level, we have chosen the regularization parameter of the Horn and Schunk method such that the error was minimal. The corresponding amplitude errors are illustrated in Fig. 4.4b. The Lucas-Kanade method is the most sensitive one, whereas the Horn and Schunk method performs reasonably well if the regularization parameter is chosen properly. The spatio-temporal-affine algorithm performs best with respect to both error measures because the multiresolution strategy and the use of a temporal window provide increased robustness against noise. However, it can be shown in general that the least-squares method, as used in (4.5), yields estimates that may be biased towards smaller values if the data matrix contains noise [62]. This can also be observed from the experiments by computing the relative bias

$$\varepsilon_b = \frac{\|\mathbf{v}\|_{l_2} - \|\widehat{\mathbf{v}}\|_{l_2}}{\|\mathbf{v}\|_{l_2}}. \quad (4.20)$$

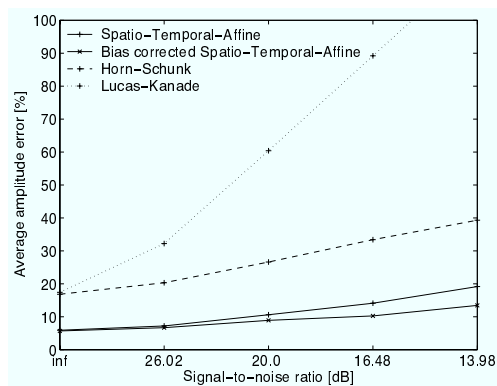
The mean bias $\bar{\varepsilon}_b$, averaged over the whole image sequence, was calculated for each noise level. The underestimation grows with increasing noise level and ranged from 1.7% to 11.4%. To demonstrate the influence of the bias to the amplitude error (4.19), we re-calculated it after correcting the velocity estimates by the average bias such that

$$\widehat{\widehat{\mathbf{v}}} = \frac{1}{1 - \bar{\varepsilon}_b} \widehat{\mathbf{v}}. \quad (4.21)$$

The resulting amplitude error is also plotted in Fig 4.4b; it is smaller than in the non-corrected case confirming our observation. A more systematic way to correct for the bias is given by the total least-squares (TLS) technique [62, 63]. However, this technique requires the knowledge of the noise covariance of a system matrix that is composed by the original data matrix and the right hand side vector. Its estimation is challenging if the sample sizes are small and if an accurate noise model is not known. Our experiments with the TLS-method are described in Chapter 5.



(a) Average angular error



(b) Average amplitude error

Figure 4.4: Average angular and amplitude errors of different methods for different signal-to-noise ratios of the point scatterer images.


The experiments demonstrate the robustness of the proposed algorithm in the case of a simple ultrasound imaging model. More realistic simulations should perhaps consider a non-linear model for the image formation and a tissue model that reflects the heart anatomy, such as fiber directions, non-uniformity of speed of sound, more closely. However, these properties are difficult to simulate which calls for alternative validation methods such as phantom experiments.

Radial Motion Visualization


For radial motion visualization, the estimated velocity was projected onto the direction to the floating centroid of the ROI to compensate for the underlying translational motion. The color-coded radial velocities for the example sequence are shown in Fig. 4.5a and 4.5b for systole and diastole, respectively. The uniform contraction during systole is clearly expressed by the red/yellow colors. On the other hand, the expansion during diastole is correctly displayed by the blue/green colors. The color display is more intuitive than the velocity field representation and is not distorted by the underlying global translation. The ROI also clearly tracks the myocardium and the circular form of the ROI-contours is maintained due to the least-squares fitting to the estimated noisy motion field.

4.4.2 Phantom Experiments

The algorithm's ability to estimate motion from real echocardiograms was tested in a phantom experiment. A cylinder-shaped, tissue-mimicking phantom was placed inside a tube of water and rotated with constant angular velocity around its longitudinal axis. The gelatine phantom had a diameter of 7.5 cm and was made out of a 2% agar-agar solution containing randomly dispersed scatterers whose size ranged from 50 to 250 μm . Echocardiograms of a 2D-cross section orthogonal to the axis of rotation were acquired and analyzed by the algorithm. Fig. 4.6a shows one frame of a rotating phantom echocardiogram and Fig. 4.6b shows its corresponding estimated velocity field. Since the applied motion is purely rotational, the velocity magnitude increases linearly with the radius. To test this relationship for the estimated velocity field, we computed local means and standard deviations of the velocity field along annuli of different radii. We then fitted a linear regression line to these mean velocities as shown in Fig. 4.7. The coefficient of determination is 0.993 which confirms the linear relationship. The slope of the regression line corresponds to the estimated angular velocity. The estimated value was 0.0727 rps, whereas the true angular velocity applied to the phantom was 0.0796 rps. This corresponds to an underestimation of 8.7%.




(a) Radial velocity during systole




(b) Radial velocity during diastole

Figure 4.5: Synthetic ultrasound sequence and superimposed color-coded radial velocity.



(a) One frame of the rotating phantom sequence



(b) Superimposed estimated velocity field

Figure 4.6: One frame of the rotating phantom sequence and its corresponding estimated velocity field.

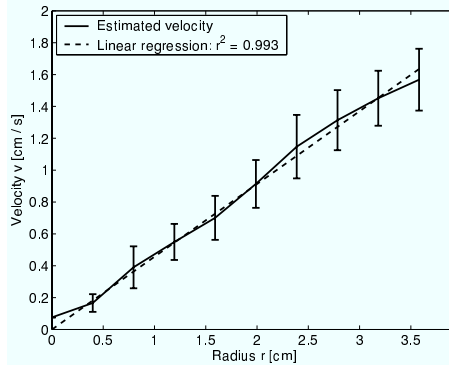



Figure 4.7: The means and standard deviations of the estimated velocity field for fixed radii are plotted together with its linear regression line.

This experiment can also serve as a calibration tool to determine the factor for the bias correction of the motion estimates as proposed in (4.21).

4.4.3 Application to Clinical Data

For a first *in vivo* validation, we applied the method to a set of clinical echocardiograms that were obtained from an animal study. We analyzed ultrasound sequences of six dogs with infarctions that were artificially induced by the occlusion of the left anterior descending artery (LAD). As introduced in Section 3.3.3 on page 29, we use the standardized division of the left ventricle into 16 segments according to the American Heart Association. Regions that depend on the left anterior descending artery blood supply are indicated in Fig. 3.6.


Echocardiograms were acquired in an open-chest state before and after the infarction. Since the infarction was caused manually, the exact localization and extend is known and a direct comparison between normal and pathological cases is possible. For one case, two frames of a long axis view (LAX) before the infarction are shown in Fig. 4.8a and 4.8c for systole and diastole, respectively. The corresponding motion estimation results are superimposed in the form of a needle diagram. The typical ventricular contraction and expansion of a normal beating heart during systole and diastole is clearly captured by the estimated motion fields. Fig. 4.9 shows the status after the infarction. Fig. 4.9a and 4.9c correspond to a frame during systole and diastole, respectively. In both figures, the dyskinesia (paradoxical motion) in the apical to mid anteroseptal segments




(a) Estimated velocity field during systole



(b) Color-coded radial velocity during systole



(c) Estimated velocity field during diastole



(d) Color-coded radial velocity during diastole

Figure 4.8: Echocardiograms with superimposed motion information during systole and diastole before infarction. The estimated velocity fields are shown in (a) and (c) and the corresponding color-coded radial velocity in (b) and (d). The velocity fields and the colors reflect the normal contraction and expansion during systole and diastole, respectively.




Figure 4.9: Echocardiograms with superimposed motion information during systole and diastole after infarction. The estimated velocity fields are shown in (a) and (c). Wall motion abnormalities (dyskinesia) are highlighted by red arrows. The corresponding color-coded radial velocity is shown in (b) and (d).

(upper right region) is correctly captured by the motion field; during systole the affected segments move outwards due to the inner blood pressure, whereas they move inwards during diastole.

The color-coded display of the radial velocity confirms the observations obtained from the motion fields, but can be interpreted in a more straightforward way. The color-coded radial velocity of the example case before the infarction is shown in Fig. 4.8. Fig. 4.8b and 4.8d show the corresponding frames during systole and diastole, respectively. In all cases, the velocity was projected onto a fixed reference point that is given by the ROI-centroid at end-diastole. The

regular contraction of the whole myocardium during systole is indicated by the red/yellow colors, whereas the blue/green colors reflect the normal expansion during diastole. The color-coded radial velocity after the infarction is shown in Fig. 4.9. In Fig. 4.9b, the outward motion of the apical to mid anteroseptal segments (upper right region) during systole is correctly visualized by the blue/green colors. The infarction also causes an outward movement of the basal posterior (lower left) segment. Fig. 4.9d shows the pathological motion during diastole. Here, the red/yellow colors indicate that the apical to mid anteroseptal segments (upper right region) move inwards instead of outwards as in the normal case. Accordingly, the basal posterior (lower left) segment also moves inwards.

Segmental Wall Motion Analysis

For all six cases, we compared peak radial velocities of myocardial segments during systole before and after the infarction. In particular, we analyzed segments that were classified as dyskinetic by the expert. Since during systole the ventricle is contracting (negative radial velocity), a normal beating heart does not exhibit significant positive radial velocities. However, dyskinetic segments expand during systole and show significant positive values. In this experiment, eight segments out of total 36 were classified as dyskinetic by the expert. The average of the measured positive peak radial velocities in these dyskinetic segments is 1.99 ± 0.92 cm/s. In contrast, the average peak velocity in these segments before the infarction was 0.18 ± 0.15 cm/s. A paired t-test shows a significant difference between the two states ($p < 0.001$). Although differences are clearly significant for the dyskinetic case, the distinction between normal, akinetic (no motion) and hypokinetic (very little motion) cases is not so well defined. Statistical evidence in these cases could not be established since the experiment only included six hearts that also show a variation in velocities before the infarction. To perform a meaningful statistical study with a higher number of infarct-categories, a larger number of cases needs to be analyzed.


ROI-Tracking Validation

The motion estimation algorithm in combination with the automatic contour tracking was also validated by comparing it to the manual contour tracking by an expert. Since the epicardial border is in general not clearly defined and is only partially imaged in our test sequences, the validation is performed by tracking the endocardial border of the left ventricle. In all 12 test sequences, an expert,


who was blind to the computer analysis, outlined the contours of the endocardial border at the end-diastolic and end-systolic frames, respectively. Each contour was defined by a set of landmarks that were interpolated by a cubic B-spline curve. In this experiment, the curves were not closed, but the two landmarks at the mitral valve attachment were connected by a straight line to define an endocardial area. The manually defined B-spline contour at end-diastole was then tracked by the algorithm to end-systole. We calibrated the system and used a constant bias correction factor in all experiments. Fig. 4.10a and Fig. 4.10c show one example of the manually determined contours (green line) at end-diastole and end-systole, respectively. The automatically tracked contour is shown in Fig. 4.10b and Fig. 4.10c for mid-systole and end-systole, respectively (white line). The areas that are enclosed by the manually determined contours at end-diastole and end-systole are denoted by the characteristic functions Ω_{ED} and Ω_{ES} , respectively. The area defined by the automatically tracked contour at end-systole is denoted as $\widehat{\Omega}_{ES}$. By calculating the area-size of a characteristic function as

$$\|\Omega\|_{L_1} = \int_{\mathbb{R}^2} |\Omega(x, y)| \, dx \, dy, \quad (4.22)$$

we define the manually determined area ejection fraction as $EF = (\|\Omega_{ED}\|_{L_1} - \|\Omega_{ES}\|_{L_1}) / \|\Omega_{ED}\|_{L_1}$. Accordingly, we denote the computer-based area ejection fraction by $\widehat{EF} = (\|\Omega_{ED}\|_{L_1} - \|\widehat{\Omega}_{ES}\|_{L_1}) / \|\Omega_{ED}\|_{L_1}$. The area ejection fraction can be interpreted as the two-dimensional analog of the volumetric ejection fraction [64] and relates the endocardial area before and after the contraction phase (systole). To compare the automatically tracked contours with the expert standard, we compute two error measures. The first is the relative ejection fraction error $E_{EF} = |EF - \widehat{EF}| / |EF|$. Let $\|\Omega_{ES} - \widehat{\Omega}_{ES}\|_{L_1}$ denote the area-size of the characteristic function that describes the region where the end-systolic manual and computer-based masks do not coincide. Then the second error measure $E_A = \|\Omega_{ES} - \widehat{\Omega}_{ES}\|_{L_1} / \|\Omega_{ES}\|_{L_1}$ corresponds to the relative size of the area, where the manual and computer-defined masks do not overlap. In this experiment, the average area error was determined as $E_A = 10.20 \pm 2.04$ % and the area ejection fraction error was $E_{EF} = 11.42 \pm 7.26$ %. To determine the variability of the expert contour tracking, the contours at end-systole were outlined a second time by the observer. The same error measures as in the manual/computer comparison were calculated. The intra-observer error for the end-systolic area was determined as $E_A = 10.56 \pm 3.45$ % and the intra-observer area ejection fraction error was $E_{EF} = 13.62 \pm 9.86$ %. These results show that the motion estimation algorithm yields realistic results and that the contours



(a) Manually defined contour at end-diastole (b) Automatically tracked contour at mid-systole



(c) Manually defined contour (green line) and automatically tracked contour (white line) at end-systole (d) Endocardial area at end-systole for which the manual and computer-based masks do not coincide

Figure 4.10: Comparison between manual and automatic contour tracking. The initial manually defined contour at end-diastole is automatically tracked to end systole and compared with the manual result.

are sufficiently well tracked to provide a time-varying ROI for further post-processing and motion visualization.

4.5 Conclusions

We proposed a new method to estimate heart motion from echocardiograms that uses a local motion model that is particularly well adapted to typical heart dynamics. The method is fast and robust due to a wavelet-like multiresolution implementation. For visualization purposes, we developed a robust method to track a ROI in time. Inside the time-varying ROI, we superimpose a color-coded radial velocity component onto the echocardiogram that allows to directly identify and quantify myocardial contractility. The algorithm outperforms other methods on simulated data; its applicability to real echocardiograms was demonstrated by a phantom experiment. A first validation of the proposed method on clinical echocardiograms yielded realistic motion fields. We also verified that the estimated regional wall motion was in good agreement with the expert echocardiographic reading.

The brain can be easy to buy, but the heart never comes to market.

— James R. Lowell

Chapter 5

Total-Least-Squares-Based Motion Estimation

Abstract — In this chapter, we consider the total-least-squares approach as an alternative to ordinary least-squares to solve the local linear systems for the motion parameters. We review the general definition of the total-least-squares problem and its solution methods. We also investigate its relation to ordinary least-squares and consider a mixture method between the two approaches—the so-called scaled total-least-squares method. We compare the feasibility of the two methods and their trade-off version on synthetic and phantom ultrasound data.

5.1 Introduction

While the least-squares method of Chapter 4 yields realistic motion estimates, one may object that this kind of method can lead to biased estimates when the matrix entries are subject to noise. In this chapter, we investigate the feasibility of an alternative approach that is based on *total-least-squares* (TLS), a general methodology for solving overdetermined linear systems, where both the matrix and the right hand side vector are subject to noise.

From statistics, it is known that the least-squares estimator as used in (4.5) on page 40 is generally biased towards lower values if the system matrix contains

noise [62]. This is best demonstrated for the Lucas-Kanade approach [19] that uses a locally constant model for the velocity in space. In this case, the Normal Equations (4.5) reduce to the (2×2) system $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$, where

$$\mathbf{A}^T \mathbf{A} = \begin{pmatrix} \langle w, I_x^2 \rangle & \langle w, I_x I_y \rangle \\ \langle w, I_x I_y \rangle & \langle w, I_y^2 \rangle \end{pmatrix}, \quad \mathbf{A}^T \mathbf{b} = - \begin{pmatrix} \langle w, I_x I_t \rangle \\ \langle w, I_y I_t \rangle \end{pmatrix}, \quad (5.1)$$

and $\mathbf{x} = (u, v)^T$ denotes the velocity components in x - and y -direction, respectively. If we assume that the spatial and temporal derivatives I_x, I_y and I_t contain additive, independently and identically distributed noise with zero-mean and variance σ^2 , the expected noise-corrupted system matrix becomes

$$\mathbf{A}^T \mathbf{A} + \sigma^2 \mathbf{I}_2 = \begin{pmatrix} \langle w, I_x^2 \rangle + \sigma^2 & \langle w, I_x I_y \rangle \\ \langle w, I_x I_y \rangle & \langle w, I_y^2 \rangle + \sigma^2 \end{pmatrix}. \quad (5.2)$$

The expectation of the right hand side vector remains $\mathbf{A}^T \mathbf{b}$. Without loss of generality, we can assume that the off-diagonal elements of (5.2) are zero because this can always be achieved by a suitable rotation of the coordinate system. Since the noise-free solution $\mathbf{x} = (x_1, x_2)^T$ satisfies $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$, the noise-corrupted solution $\tilde{\mathbf{x}}$ can be expressed in terms of \mathbf{x} as

$$\tilde{\mathbf{x}} = (\mathbf{A}^T \mathbf{A} + \sigma^2 \mathbf{I}_2)^{-1} \mathbf{A}^T \mathbf{b} = \begin{pmatrix} \frac{\langle w, I_x^2 \rangle}{\langle w, I_x^2 \rangle + \sigma^2} x_1 \\ \frac{\langle w, I_y^2 \rangle}{\langle w, I_y^2 \rangle + \sigma^2} x_2 \end{pmatrix}. \quad (5.3)$$

Equation (5.3) demonstrates that the solution in the noisy case is biased towards lower values. Furthermore, the bias also depends on the signal-to-noise ratio of the spatial derivatives.

This chapter is organized as follows. We will first review the general definition of the total-least-squares problem in Section 5.2. Its solution and its relation to ordinary least-squares are investigated in Sections 5.2.1 and 5.2.2, respectively. In Section 5.2.3, we consider a mixture method between ordinary and total-least-squares, known as scaled total-least-squares. This method allows to continuously switch between the two methods by varying a single trade-off parameter. We test the applicability of these approaches on synthetic and phantom ultrasound data in Section 5.3.

5.2 The Total-Least-Squares Problem

The total-least-squares method (TLS), also known as errors-in-variables method in statistics, has been proposed to take into account the uncertainties in the data matrix \mathbf{A} . In the following, we give a brief summary of the total-least-squares approach; an extended analysis can be found in [62, 63, 65].

In the classical least-squares formulation, it is assumed that the data errors are confined to the right hand side vector \mathbf{b} . However, this assumption is not true in our case since the data matrix \mathbf{A} is composed of the spatial derivatives that are plagued with uncertainties due to the noise in the images. Let $\mathbf{Ax} \approx \mathbf{b}$ be an overdetermined, full rank linear system of equations in the unknowns $\mathbf{x} \in \mathbb{R}^n$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, $m > n$, and $\mathbf{b} \notin \text{Image}(\mathbf{A})$. The ordinary least-squares formulation can be interpreted as finding a solution \mathbf{x} and a vector $\hat{\mathbf{b}}$ such that:

$$\min_{\hat{\mathbf{b}}} \|\hat{\mathbf{b}} - \mathbf{b}\|_{l_2} \quad (5.4)$$

subject to

$$\mathbf{Ax} = \hat{\mathbf{b}}. \quad (5.5)$$

By using standard algebraic manipulations, we find that the minimum of the least-squares problem is equal to the solution of the so-called Normal Equations; i.e.,

$$\mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}. \quad (5.6)$$

The idea behind total-least-squares is to consider perturbations of both \mathbf{b} and \mathbf{A} . More precisely, the total-least-squares problem consists in finding a solution \mathbf{x} , a matrix $\hat{\mathbf{A}}$ and a vector $\hat{\mathbf{b}}$ such that:

$$\min_{\hat{\mathbf{A}}, \hat{\mathbf{b}}} \|(\hat{\mathbf{A}}|\hat{\mathbf{b}}) - (\mathbf{A}|\mathbf{b})\|_F \quad (5.7)$$

subject to

$$\hat{\mathbf{A}}\mathbf{x} = \hat{\mathbf{b}}. \quad (5.8)$$

Here, $\|\cdot\|_F$ denotes the Frobenius norm and the matrix $(\mathbf{A}|\mathbf{b}) \in \mathbb{R}^{m \times (n+1)}$ denotes the column-wise extended matrix \mathbf{A} by \mathbf{b} . If the rows of the error matrix $\mathbf{E} = (\hat{\mathbf{A}}|\hat{\mathbf{b}}) - (\mathbf{A}|\mathbf{b})$ are independently and identically distributed with zero mean

vector and common covariance matrix $\mathbf{C} = \sigma^2 \mathbf{I}_{n+1}$ with some variance σ^2 , it can be shown that the total-least-squares solution corresponds to the maximum likelihood estimator of \mathbf{x} [62].

As in the least-squares approach, the influence of each equation in the system $\mathbf{A}\mathbf{x} \approx \mathbf{b}$ to the solution can be adjusted by multiplying the system by a left hand side weighting matrix $\mathbf{W} \in \mathbb{R}^{m \times m}$, which is typically a diagonal matrix, where each entry expresses our confidence in the corresponding row of the linear system. In the following, we assume that \mathbf{A} and \mathbf{b} already correspond to the weighted versions.

5.2.1 Solution of the Total-Least-Squares Problem

By rewriting $\mathbf{A}\mathbf{x} \approx \mathbf{b}$ in the form

$$(\mathbf{A}|\mathbf{b}) \begin{pmatrix} \mathbf{x} \\ -1 \end{pmatrix} \approx \mathbf{0}, \quad (5.9)$$

it can be seen that the total-least-squares problem corresponds to finding a perturbation matrix $\mathbf{E} = (\widehat{\mathbf{A}}|\widehat{\mathbf{b}}) - (\mathbf{A}|\mathbf{b})$ of minimal Frobenius norm such that $(\mathbf{A}|\mathbf{b}) + \mathbf{E}$ is rank deficient. Let the singular value decomposition of $(\mathbf{A}|\mathbf{b})$ be given by

$$\mathbf{U}^T (\mathbf{A}|\mathbf{b}) \mathbf{V} = \mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_{n+1}) \quad (5.10)$$

with orthogonal matrices $\mathbf{U} \in \mathbb{R}^{m \times m}$, $\mathbf{V} \in \mathbb{R}^{(n+1) \times (n+1)}$ and singular values $\sigma_1 \geq \dots \geq \sigma_i \geq \dots \geq \sigma_{n+1}$. Then, it can be shown that

$$\sigma_{n+1} = \min_{\text{rank}((\widehat{\mathbf{A}}|\widehat{\mathbf{b}}))=n} \|(\mathbf{A}|\mathbf{b}) - (\widehat{\mathbf{A}}|\widehat{\mathbf{b}})\|_F. \quad (5.11)$$

Moreover, if $\sigma_n > \sigma_{n+1}$ and $\mathbf{V}_{n+1,n+1} \neq 0$, then

$$\mathbf{x} = -\frac{1}{\mathbf{V}_{n+1,n+1}} (\mathbf{V}_{1,n+1}, \dots, \mathbf{V}_{n,n+1})^T \quad (5.12)$$

exists and is the unique solution of $\mathbf{A}\mathbf{x} \approx \mathbf{b}$. Thus, the total-least-squares solution is given by the first n components of the right singular vector that belongs to the minimal singular value σ_{n+1} of $(\mathbf{A}|\mathbf{b})$. The singular vector is normalized by its negative last component to satisfy formulation (5.9).

5.2.2 Relation between Least-Squares and Total-Least-Squares

By simple algebraic manipulations, it can be shown that the right hand side singular vectors and squared singular values σ_i^2 of $(\mathbf{A}|\mathbf{b})$ correspond to the eigenvectors and eigenvalues of $(\mathbf{A}|\mathbf{b})^T(\mathbf{A}|\mathbf{b})$, respectively. Since, $(\mathbf{x}, -1)^T$ is a singular vector with singular value σ_{n+1} of $(\mathbf{A}|\mathbf{b})$, it satisfies the following eigenvalue relation:

$$\begin{aligned} (\mathbf{A}|\mathbf{b})^T(\mathbf{A}|\mathbf{b}) \begin{pmatrix} \mathbf{x} \\ -1 \end{pmatrix} = \\ \begin{pmatrix} \mathbf{A}^T\mathbf{A} & \mathbf{A}^T\mathbf{b} \\ \mathbf{b}^T\mathbf{A} & \mathbf{b}^T\mathbf{b} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ -1 \end{pmatrix} = \sigma_{n+1}^2 \begin{pmatrix} \mathbf{x} \\ -1 \end{pmatrix}. \end{aligned} \quad (5.13)$$

From the upper part of the block-structured system (5.13), it follows immediately that

$$\mathbf{x} = (\mathbf{A}^T\mathbf{A} - \sigma_{n+1}^2\mathbf{I}_{n+1})^{-1}\mathbf{A}^T\mathbf{b}. \quad (5.14)$$

Equation (5.14) indicates that the total-least-squares formulation takes into account the additional bias term that is due to noise as exemplified in (5.3). Under suitable conditions, it can be shown that the smallest squared singular value σ_{n+1}^2 is a consistent estimator of the true noise variance σ^2 [62], i.e., $\lim_{m \rightarrow \infty} \sigma_{n+1}^2/m = \sigma^2$. Since the total-least-squares formulation consists in subtracting a multiple of the identity matrix, it can be interpreted as a de-regularization method, in contrast to the so-called regularized least-squares estimators, which add a positive definite term to the system matrix. Therefore, the total-least-squares formulation is usually even less well conditioned than the least-squared method.

Equation (5.13) reveals also an alternative way to determine the total-least-squares solution. Since for motion estimation, we directly compute the Normal Equations in terms of $\mathbf{A}^T\mathbf{A}$ and $\mathbf{A}^T\mathbf{b}$, the terms \mathbf{A} and \mathbf{b} are not known explicitly. Therefore, we use the least-squares terms to assemble the block-structured system matrix $(\mathbf{A}|\mathbf{b})^T(\mathbf{A}|\mathbf{b})$ as shown in (5.13). The only term that needs to be computed additionally is the scalar product $\mathbf{b}^T\mathbf{b}$. Computing the eigenvalues and eigenvectors of this system matrix yields the desired total-least-squares solution. Since this $(n+1) \times (n+1)$ matrix is usually much smaller than the original measurement matrix $(\mathbf{A}|\mathbf{b}) \in \mathbb{R}^{m \times (n+1)}$, its eigenvalues and eigenvectors can be computed efficiently.

The geometrical interpretation of the total-least-squares approach for motion estimation can be most easily demonstrated with the Lucas-Kanade approach.

If we consider the image sequence as a three-dimensional volume, the direction of motion can be interpreted as the unit vector $\mathbf{u} = (u_1, u_2, u_3)^T$ that points into the direction of minimal intensity variation; i.e., the vector that is orthogonal to the spatio-temporal gradient $\nabla I = (I_x, I_y, I_t)^T$. Considering a local spatio-temporal neighborhood around some position (x_0, y_0, t_0) and some weighting function w , the unit vector \mathbf{u} can be determined by minimizing the following weighted least-squares criterion [66]:

$$\min_{\mathbf{u}} \int_{\mathbb{R}^3} w(x - x_0, y - y_0, t - t_0) \left(\nabla I^T \cdot \mathbf{u} \right)^2 dx dy dt \quad (5.15)$$

subject to

$$\|\mathbf{u}\|_2 = 1.$$

By rewriting $(\nabla I^T \cdot \mathbf{u})^2 = \mathbf{u}^T \nabla I \nabla I^T \mathbf{u}$ and introducing the Lagrange multiplier $\lambda \in \mathbb{R}$, the constrained optimization (5.15) is equivalently expressed as:

$$\min_{\mathbf{u}, \lambda} \mathbf{u}^T \mathbf{G} \mathbf{u} + \lambda(1 - \|\mathbf{u}\|_2), \quad (5.16)$$

where

$$\mathbf{G} = \begin{pmatrix} \langle w, I_x^2 \rangle & \langle w, I_x I_y \rangle & \langle w, I_x I_t \rangle \\ \langle w, I_x I_y \rangle & \langle w, I_y^2 \rangle & \langle w, I_y I_t \rangle \\ \langle w, I_x I_t \rangle & \langle w, I_y I_t \rangle & \langle w, I_t^2 \rangle \end{pmatrix}. \quad (5.17)$$

This leads to the eigenvalue problem

$$\mathbf{G} \mathbf{u} = \lambda \mathbf{u}. \quad (5.18)$$

The minimum is obtained if \mathbf{u} corresponds to the eigenvector to the minimal eigenvalue λ of \mathbf{G} . As can be seen from (5.1) and (5.13), the matrix \mathbf{G} is nothing but the system matrix $(\mathbf{A}|\mathbf{b})^T(\mathbf{A}|\mathbf{b})$ of the total-least-squares formulation of the Lucas-Kanade optical-flow approach. Interestingly enough, this total-least-squares system is equivalent to the *structure tensor* approach for multidimensional orientation estimation [67].

5.2.3 Scaled Total-Least-Squares

The main weakness of the total-least-squares method is that it may overcorrect the matrix \mathbf{A} by replacing it by some matrix $\hat{\mathbf{A}}$ that is too far from it. This

can be prevented by imposing some *a priori* bound on the size of the allowable correction [68, 69]. Unfortunately, this can only be achieved by means of iterative algorithms that require a much higher computational load than the least-squares and total-least-squares approaches. In addition, they only converge under certain conditions that are not always satisfied. Another way to control the potential overcorrection is by using a mixture formulation of ordinary and total-least-squares. This can be implemented with the so-called scaled total-least-squares method [70] which seeks a matrix $\widehat{\mathbf{A}}$ and vectors $\widehat{\mathbf{b}}$ and \mathbf{x} such that for a given real number $\gamma \in]0, 1]$ the following functional is optimized:

$$\min_{\widehat{\mathbf{A}}, \widehat{\mathbf{b}}} \|(\widehat{\mathbf{A}}|\gamma\widehat{\mathbf{b}}) - (\mathbf{A}|\gamma\mathbf{b})\|_F \quad (5.19)$$

subject to

$$\widehat{\mathbf{A}}\mathbf{x} = \widehat{\mathbf{b}}. \quad (5.20)$$

For $\gamma = 1$, this method is obviously equivalent to standard total-least-squares. In fact, the scaled version can be implemented by using the standard algorithm and simply replacing \mathbf{b} by the rescaled vector $\gamma\mathbf{b}$. Moreover, by using (5.14), we can show that the scaled total-least-squares method converges to the ordinary least-squares solution for $\gamma \rightarrow 0$. For further details, we refer to [70]. Thus, for a suitable choice of $\gamma \in]0, 1]$, we can achieve a compromise between the ordinary and total-least-squares approaches, or in other words, a trade-off between decreasing the bias while increasing the variance of the estimator.

5.2.4 Equilibration of the Total-Least-Squares System

In the considered case of motion estimation, the entries of the measurement matrix \mathbf{A} and right hand side vector \mathbf{b} are likely to be correlated and their variance may not be identical. A consistent estimate of the model parameters can be obtained by pre-whitening the augmented system matrix $(\mathbf{A}|\mathbf{b})$ such that the expectation $E(\mathbf{E}^T\mathbf{E})$ becomes diagonal with equal error variances [71]. This can be achieved by applying a suitably chosen, non-singular right hand side equilibration matrix $\mathbf{R} \in \mathbb{R}^{(n+1) \times (n+1)}$ that satisfies

$$E(\mathbf{R}^T\mathbf{E}^T\mathbf{E}\mathbf{R}) = \sigma^2\mathbf{I}_{n+1} \quad (5.21)$$

for some noise variance σ^2 . Since $E(\mathbf{R}^T\mathbf{E}^T\mathbf{E}\mathbf{R}) = \mathbf{R}^TE(\mathbf{E}^T\mathbf{E})\mathbf{R}$, the equilibration matrix \mathbf{R} is given by the root of the inverse covariance matrix, i.e.,

$\mathbf{R}\mathbf{R}^T = E(\mathbf{E}^T\mathbf{E})^{-1}$. The equilibrated total-least-squares solution of the linear system $\mathbf{A}\mathbf{x} \approx \mathbf{b}$ is then obtained by first determining the total-least-squares solution of

$$((\mathbf{A}|\mathbf{b})\mathbf{R})\mathbf{y} \approx \mathbf{0} \quad (5.22)$$

as described in Section 5.2.1. Here, $\mathbf{y} = \mathbf{R}^{-1}(\mathbf{x}, -1)^T \in \mathbb{R}^{n+1}$ denotes the intermediate solution of the equilibrated system; the final solution is given by $\mathbf{x} = -1/z_{n+1} \cdot (z_1, \dots, z_n)^T$, where $\mathbf{z} = \mathbf{R}\mathbf{y} \in \mathbb{R}^{n+1}$.

5.3 Numerical Results

The applicability of the total-least-squares approach and its scaled variation was tested in combination with the spatio-temporal-affine motion model introduced in Section 4.2.1. The equilibration matrix needed to whiten the system matrix, assembled from the Normal Equations (4.5) on page 40, is derived as follows: if we assume that the spatial and temporal derivatives I_x, I_y and I_t contain additive, zero-mean, identically, and independently noise of common variance σ^2 , we find that

$$E(\mathbf{E}^T\mathbf{E}) = \sigma^2 \text{diag}\left(\|w\|_1, \|w\|_1, \|x^2w\|_1, \|y^2w\|_1, \|t^2w\|_1, \|x^2w\|_1, \|y^2w\|_1, \|t^2w\|_1, \|w\|_1\right), \quad (5.23)$$

where w denotes the discretized window function. The off-diagonal entries are zero due to the independency assumption of the noise. The corresponding equilibration matrix \mathbf{R} is given by the inverse component-wise square root of the diagonal matrix $E(\mathbf{E}^T\mathbf{E})$, i.e.,

$$\mathbf{R} = \frac{1}{\sigma} \text{diag}\left(\|w\|_1^{-1/2}, \|w\|_1^{-1/2}, \|x^2w\|_1^{-1/2}, \|y^2w\|_1^{-1/2}, \|t^2w\|_1^{-1/2}, \|x^2w\|_1^{-1/2}, \|y^2w\|_1^{-1/2}, \|t^2w\|_1^{-1/2}, \|w\|_1^{-1/2}\right). \quad (5.24)$$

We computed the corresponding equilibration matrix for each dilated version $w^{(j)}$ of the window used. For the coarse-to-fine multi-scale refinement scheme described in Section 4.2.3, we used a three-level pyramid decomposition during our experiments.

Application to Synthetic Ultrasound Data

To compare the performance of the different approaches, we applied them first to the synthetic ultrasound sequences described in Section 4.4.1. In particular, we assessed the robustness of the methods by using echo sequences that were obtained from point scatterer images of signal-to-noise ratios varying from 26.02 to 6.01 dB. The maximum applied velocity in these sequences was 2.0 pixel/frame.

The difference between the estimated velocity $\hat{\mathbf{v}}$ and the exact velocity \mathbf{v} is measured as


$$\varepsilon = \|\hat{\mathbf{v}}\|_{l_2} - \|\mathbf{v}\|_{l_2}. \quad (5.25)$$

The average bias $\bar{\varepsilon}$ was defined as the mean of the velocity differences ε computed over all pixels in the image sequence with non-zero motion. The standard deviation σ_ε of ε was computed analogously.

Fig. 5.1 shows the results obtained for the ordinary least-squares method ($\gamma \approx 0$), the total-least-squares method ($\gamma = 1$), and the scaled total-least-squares method ($\gamma = 0.5$). The total-least-squares method only produced velocity estimates in an acceptable range for a fraction of the image data. As mentioned in Section 4.2.3, motion vectors are rejected if their size exceeds a scale-dependent limit. As plotted in Fig. 5.1c, the motion field density of the total-least-squares method decreases for increasing noise level. The ordinary least-squares and scaled total-least-squares method yielded realistic results everywhere. Fig. 5.1a demonstrates that the ordinary least-squares method shows an increasing bias towards smaller velocity magnitudes as the noise level increases. On the other hand, the total-least-squares method tends to overcorrect this bias, leading to an overestimate of the velocities. The combination of both methods produces the smallest bias. As expected, the total-least-squares method shows the largest variance of the motion estimates, plotted in Fig. 5.1b. Again, the scaled total-least-squares method performs best.


Application to Phantom Ultrasound Data

The same experiment was also performed on the rotating phantom ultrasound data described in Section 4.4.2. The maximum velocity magnitude of the sequence used here was 6.6 pixel/frame. Analogously to the experiments on synthetic data, the least-squares method produced a bias of -0.59 pixels/frame towards lower velocities. The total-least-squares method overestimated the velocities on average by 0.94 pixel/frame. The scaled total-least-squares method



(a) Means of differences between estimated and correct velocities

(b) Standard deviations of differences between estimated and correct velocities



(c) Motion field density of total-least-squares method

Figure 5.1: Performance comparison of ordinary least-squares (LS), total-least-squares (TLS) and scaled total-least-squares on synthetic ultrasound data. Whereas the LS-method underestimates the velocity magnitudes, the TLS-method tends to overestimate them (a). The scaled TLS-approach yields the smallest bias (a) and standard deviation (b). The applicability of the TLS-approach reduces with increasing noise level in the ultrasound data (c).

($\gamma = 0.5$) yielded the smallest bias of 0.11 pixel/frame. It also showed the smallest standard deviation of the velocity differences.

5.4 Conclusions

We have tested the feasibility of the total-least-squares method as an alternative approach to fit the local motion parameters. Given the amount of noise present in the ultrasound data, this method tends to be too sensitive to be applied routinely. The experimental results suggest that the scaled total-least-squares method is a good compromise between the two approaches. The performance, however, depends on the trade-off parameter which needs to be carefully chosen. To render the application of total-least-squares more robust, one has to employ appropriate confidence or coherence measures that allow to detect cases of non-applicability. Since the ordinary least-squares method yields robust results on a wide range of ultrasound data, we use this method throughout our experiments if not mentioned otherwise.

Two things are bad for the heart—running up stairs and running down people.

— Bernard M. Baruch

Chapter 6

Bimodal Motion Analysis from B-mode and Tissue Doppler Ultrasound

Abstract — We present a new method for estimating heart motion from two-dimensional echocardiographic sequences by exploiting two ultrasound modalities: B-mode and tissue Doppler. The algorithm estimates a two-dimensional velocity field locally by using a spatial affine velocity model inside a sliding window. Within each window, we minimize a local cost function that is composed of two quadratic terms: an optical-flow constraint that involves the B-mode data and a constraint that enforces the agreement of the velocity field with the directional tissue Doppler measurements. The relative influence of the two different modalities to the resulting solution is controlled by an adjustable weighting parameter. Robustness is achieved by a coarse-to-fine multi-scale approach. The method was tested on synthetic ultrasound data and validated by a rotating phantom experiment. First applications to clinical echocardiograms give promising results.

This chapter is adapted from our paper [72].

6.1 Introduction

Echocardiography is an established clinical tool for the assessment of regional ventricular function in patients with cardiac disease. B-mode and tissue Doppler modalities, in particular, are widely used to analyze myocardial motion in order to evaluate the degree of ischemia and infarction. B-mode echocardiograms provide a two-dimensional gray scale display of the moving heart but no quantitative motion information. Therefore, the assessment of myocardial motion from B-mode echocardiograms is routinely accomplished by visual interpretation and manual evaluation. This leads to subjective and, at best, semi-quantitative diagnoses, which suffer from a significant intra- and inter-observer variability.

Tissue Doppler imaging (TDI) is an emerging ultrasound modality that has been introduced recently as a modification of Doppler-based blood flow measurement [73, 74, 75, 76, 77]. It analyzes high-amplitude, low-frequency Doppler shifts of reflected echoes to quantify tissue motion. Among others, TDI has been used to detect wall motion abnormalities in stress echocardiography [49]. Low Doppler velocities have proven to be a strong indicator for ischemia [78, 79]. However, TDI measures only the velocity component along the ultrasonic beam direction. The fact that ventricular motion is multi-dimensional and usually not restricted to the scan line makes its interpretation difficult—especially since the images are scanned in a radial fashion.

To obtain true two-dimensional, quantitative motion information, different methods have been proposed to estimate ventricular motion from B-mode echocardiograms [17, 21, 48]. In this chapter, we propose a new motion analysis approach that is based on both modalities, B-mode and tissue Doppler. The idea is to estimate a velocity field from B-mode echocardiograms such that its projections along the scan lines are in good agreement with the tissue Doppler measurements.

The chapter is organized as follows. In Section 6.2, we briefly review the principles of tissue Doppler imaging. After reviewing the optical-flow principle to analyze motion from B-mode echocardiograms in Section 6.3, we present the proposed bimodal algorithm in Section 6.4. Numerical results from synthetic data, phantom experiments and clinical echocardiograms are given in Section 6.5.

6.2 Principles of Tissue Doppler Imaging

Velocity measurements using pulsed-wave ultrasound have become an important diagnostic tool in echocardiography [79, 49]. In two-dimensional tissue Doppler imaging (TDI), the beam is scanned radially over the region to be imaged. As illustrated in Fig. 6.1, the system measures the tissue velocity component

$$v_D = \|\mathbf{v}\|_2 \cos(\beta)$$

at positions along each beam, where β denotes the angle between the true velocity vector \mathbf{v} and the beam direction α . Thus, the measured motion depends highly on β . In particular, the modality is totally blind to displacements orthogonal to the scan line.

Conventional Doppler ultrasonography is based on the principle that v_D is proportional to the small frequency shift f_D in the ultrasound carrier frequency f_0 between transmitted and received echoes, i.e.,

$$v_D = \frac{c}{2f_0} f_D,$$

where c denotes the speed of sound. Because of the additional frequency downshift due to attenuation and the requirement of small bandwidth pulses (long duration), other methods, such as measuring deviations in time delays between successively received pulses [45], have been proposed to determine v_D .

Current ultrasound systems allow the simultaneous acquisition of B-mode and tissue Doppler signals in real-time. The Doppler velocities are usually color-coded and superimposed onto the B-mode echocardiogram as shown in Fig. 6.5b on page 86. However, the interpretation of these radial projections of complex motion patterns, such as translation, rotation and shear, is difficult and requires a sufficient level of experience [49].

6.3 Motion Analysis from B-mode Echocardiograms

Since tissue Doppler imaging is limited to the ultrasonic beam direction, several methods have been proposed to estimate true two-dimensional motion from dynamic B-mode echocardiograms; optical-flow methods, in particular, have led to promising results [21]. Gradient-based optical-flow estimation relies on the assumption that the intensity of a particular point in a moving pattern does




Figure 6.1: The tissue Doppler velocity \mathbf{v}_D corresponds to the projection of the true velocity \mathbf{v} along the beam direction α .

not change with time. Let $I(x, y, t)$ denote the intensity of pixels at location $\mathbf{r} = (x, y)$ and time t in an image sequence. Then the constant intensity assumption can be expressed as [16]

$$I_x(\mathbf{r}, t)u(\mathbf{r}, t) + I_y(\mathbf{r}, t)v(\mathbf{r}, t) + I_t(\mathbf{r}, t) = 0, \quad (6.1)$$

where I_x , I_y and I_t denote the spatial and temporal derivatives of the image intensity. The velocities u and v are, respectively, the x - and y -components of the optical-flow that we wish to estimate. Since (6.1) is a single equation in two unknowns u and v , it cannot be solved uniquely without introducing additional constraints. The Lucas-Kanade method [21], for instance, assumes the velocity to be constant within small spatial neighborhoods. It has been applied to echocardiograms by Chunke *et al.* [20].

6.4 Bimodal Motion Analysis

Since both modalities, B-mode and tissue Doppler, provide valuable motion information, we propose a novel algorithm that integrates these two kinds of information to estimate a true two-dimensional velocity field. Inspired by the Lucas-Kanade method, we propose a sliding-window algorithm. Since typical heart motions are given by rotation, expansion, contraction, and shear, we use a local affine model for the velocity in space. Let $\mathbf{r}_0 = (x_0, y_0)^T$ denote the center

of a small image region for a frame at time t . Omitting the temporal parameter for notational convenience, the spatial affine model is defined as

$$\mathbf{v}(\mathbf{r}) = \mathbf{v}_0 + \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix} \cdot (\mathbf{r} - \mathbf{r}_0), \quad (6.2)$$

where $\mathbf{r} = (x, y)^T$ and $\mathbf{v}(\mathbf{r}) = (u(\mathbf{r}), v(\mathbf{r}))^T$. The vector \mathbf{v}_0 corresponds to the velocity at the center point \mathbf{r}_0 ; u_x, u_y, v_x , and v_y are the first order spatial derivatives of u and v , respectively. These derivatives are assumed to be constant within the local neighborhood.

We estimate the local model parameters by minimizing the weighted least-squares criterion

$$\begin{aligned} & \sum_n w(\mathbf{r}_n - \mathbf{r}_0) \left((I_x(\mathbf{r}_n) u(\mathbf{r}_n) + I_y(\mathbf{r}_n) v(\mathbf{r}_n) + I_t(\mathbf{r}_n))^2 \right. \\ & \left. + \lambda (\cos(\alpha_n) u(\mathbf{r}_n) + \sin(\alpha_n) v(\mathbf{r}_n) - v_D(\mathbf{r}_n))^2 \right), \end{aligned} \quad (6.3)$$

where the sum is taken over all pixels \mathbf{r}_n inside an observation window w centered at position \mathbf{r}_0 . The first term of the cost function is equal to the optical-flow constraint applied to the B-mode data. The second term corresponds to the difference between the measured tissue Doppler velocities and the velocity field projection along the scan lines. The scan line direction at position \mathbf{r}_n is given by the unit vector $(\cos(\alpha_n), \sin(\alpha_n))$. The relative influence of the B-mode data versus the tissue Doppler measurements to the resulting velocity estimate is controlled by the non-negative trade-off parameter λ . The symmetric window function w gives more weight to constraints at the center of the local spatial region than to those at the periphery. A well suited window function is $w(x, y) = \beta^n(x)\beta^n(y)$, where β^n is the symmetrical B-spline of degree $n \in \mathbb{N}$ [51]. B-splines rapidly converge to Gaussians when their degree increases which ensures isotropy of the window in multiple dimensions.

By inserting (6.2) into (6.3) and differentiating the latter one with respect to each of the six unknown model parameters $\mathbf{x} = (u_0, v_0, u_x, u_y, v_x, v_y)^T$, we obtain the symmetric linear system (6.5) on page 82 that has to be solved at each window position. The coefficients of these Normal Equations $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$ at a given window position \mathbf{r}_0 have the form of the local moments

$$m_{p,q}(x_0, y_0) = \sum_n (x_n - x_0)^p (y_n - y_0)^q w(x_n - x_0, y_n - y_0) F_{a,b}(x_n, y_n), \quad (6.4)$$

where $0 \leq p + q \leq 2$. Since the optical-flow constraint (i.e., $I_x u + I_y v + I_t = 0$) and the tissue Doppler constraint (i.e., $\cos(\alpha)u + \sin(\alpha)v - v_D = 0$) have indeed

$$\begin{aligned}
& \mathbf{A}^T \mathbf{A} = \\
& \begin{pmatrix} \langle w, F_{x,x} \rangle & \langle w, F_{x,y} \rangle & \langle xw, F_{x,x} \rangle & \langle yw, F_{x,x} \rangle & \langle xw, F_{x,y} \rangle & \langle yw, F_{x,y} \rangle \\ \cdot & \langle w, F_{y,y} \rangle & \langle xw, F_{x,y} \rangle & \langle yw, F_{x,y} \rangle & \langle xw, F_{y,y} \rangle & \langle yw, F_{y,y} \rangle \\ \cdot & \cdot & \langle x^2 w, F_{x,x} \rangle & \langle xyw, F_{x,x} \rangle & \langle x^2 w, F_{x,y} \rangle & \langle xyw, F_{x,y} \rangle \\ \cdot & \cdot & \cdot & \langle y^2 w, F_{x,x} \rangle & \langle xyw, F_{x,y} \rangle & \langle y^2 w, F_{x,y} \rangle \\ \cdot & \cdot & \cdot & \cdot & \langle x^2 w, F_{y,y} \rangle & \langle xyw, F_{y,y} \rangle \\ \cdot & \cdot & \cdot & \cdot & \cdot & \langle y^2 w, F_{y,y} \rangle \end{pmatrix} \\
& \mathbf{x} = \begin{pmatrix} u_0 \\ v_0 \\ u_x \\ u_y \\ v_x \\ v_y \end{pmatrix}, \quad \mathbf{A}^T \mathbf{b} = - \begin{pmatrix} \langle w, F_{x,t} \rangle \\ \langle w, F_{y,t} \rangle \\ \langle xw, F_{x,t} \rangle \\ \langle yw, F_{x,t} \rangle \\ \langle xw, F_{y,t} \rangle \\ \langle yw, F_{y,t} \rangle \end{pmatrix} \quad (6.5)
\end{aligned}$$

the same structure, we first add up corresponding terms to functions denoted by $F_{a,b}$. In particular, to unify notation, we denote the unit vector along the scan line direction α at position \mathbf{r} as $(p_x(\mathbf{r}), p_y(\mathbf{r})) = (\cos(\alpha(\mathbf{r})), \sin(\alpha(\mathbf{r})))$ and define $p_t(\mathbf{r}) = -v_D(\mathbf{r})$. Note that the subscripts of p do not denote partial derivatives here. Using this notation, the functions $F_{a,b}$ are defined as

$$F_{a,b}(\mathbf{r}) = I_a(\mathbf{r}) I_b(\mathbf{r}) + \lambda p_a(\mathbf{r}) p_b(\mathbf{r}),$$

where a and b correspond to one of the variables x, y or t , respectively. From these functions, we compute the coefficients (6.4) by applying the corresponding moment filters.

6.4.1 Coarse-To-Fine Multi-Scale Strategy

To be able to estimate large motions and to increase the robustness to noise, we analogously apply the coarse-to-fine multi-scale strategy described in Section 4.2.3 on page 43. We compute a least-squares pyramid for each frame of the B-mode image sequence $I(x, y, t)$. The partial derivatives $I_x^{(k)}$, $I_y^{(k)}$ and $I_t^{(k)}$, at level k , are obtained from the corresponding B-spline representations of $I^{(k)}$. The tissue Doppler velocities $p_t = -v_D$ are treated in the same way as the B-mode data I by computing the corresponding least-squares pyramids $p_t^{(k)}$ for each frame. At each pyramid level k , the Doppler velocities are divided by 2 to account for the downsampling. The Doppler-projection operators $p_x = \cos(\alpha)$ and $p_y = \sin(\alpha)$, on the other hand, are directly computed at each resolution

level from the corresponding grid. Since they do not depend on time, they need to be computed only once for each level.

The size of the local window functions is chosen adaptively to the underlying motion field as described in Section 4.2.2 on page 39.

6.5 Numerical Results

For validation purposes, the algorithm was tested on synthetic data, on ultrasound sequences of a rotating phantom, and on clinical echocardiograms. All echocardiograms were acquired with a HP Sonos 5500 ultrasound system.

6.5.1 Application to Synthetic Data

The algorithm was first tested on synthetic ultrasound sequences for which the exact motion field and corresponding tissue Doppler velocities are known. The data was generated by warping a single reference frame of point scatterers. To simulate noise, we added scatterers of random echogenicity to each frame. The final B-mode images were obtained by applying a simple, linear ultrasound imaging model to the perturbed point scatterer images. The example sequence used here simulates a full cardiac cycle of a left ventricular short axis view (SAX) with an underlying translation to the upper right. Fig. 6.2a shows one diastolic frame of the sequence together with the superimposed estimated velocity field. The color-coded simulated tissue Doppler velocities are shown in Fig. 6.2b.

To assess the performance of the algorithm, we use the angular error measure

$$\theta = \arccos \left(\frac{\langle \mathbf{v}, \hat{\mathbf{v}} \rangle}{\|\mathbf{v}\|_2 \|\hat{\mathbf{v}}\|_2} \right).$$

between the estimated velocity $\hat{\mathbf{v}}$ and the exact velocity \mathbf{v} . The mean angular error $\bar{\theta}$ is computed by averaging θ over the whole image sequence. Fig. 6.4a illustrates the average error for different values of the trade-off parameter λ that controls the relative influence of B-mode versus tissue Doppler data in (6.3). First, the mean error decreases significantly with increasing weight of the tissue Doppler term. The relative error improvement between $\lambda = 0$ (only B-mode data is used) and $\lambda = 500$ is 19.82%. For larger values of λ , the error increases again and the local linear systems become more and more ill-conditioned.

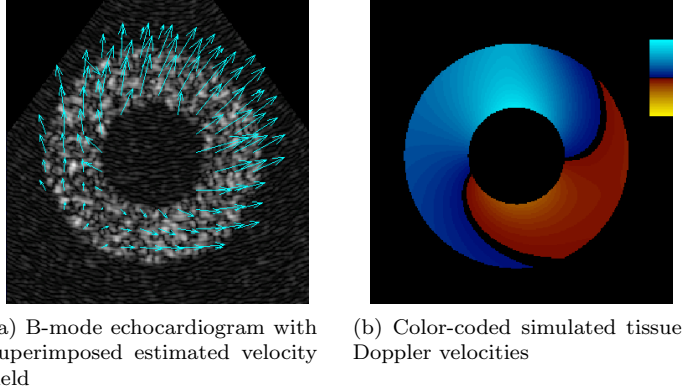


Figure 6.2: Frame of synthetic echocardiogram during diastole (a) and simulated Doppler signal (b).

6.5.2 Phantom Experiments

The algorithm's ability to analyze motion from real ultrasound data was tested by a phantom experiment. A cylinder-shaped, tissue-mimicking phantom was placed inside a tube of water and rotated with constant angular velocity around its longitudinal axis. Fig. 6.3a shows one frame of the B-mode sequence together with the superimposed estimated velocity field. Fig. 6.3b shows the measured tissue Doppler velocities. The central vertical line of zero Doppler velocities clearly demonstrates that motion orthogonal to the beam direction cannot be measured by TDI, whereas the proposed algorithm is independent from the beam direction. The average angular error θ in dependency of λ is plotted in Fig. 6.4b. The optimal error improvement of 18.15% is obtained for $\lambda = 350$. As in the case of synthetic data, the error increases if either of the two modalities becomes dominant.

6.5.3 Application to Clinical Data

For a first *in vivo* validation, we applied the method to a set of clinical echocardiograms. Fig. 6.5a shows one frame of a B-mode sequence during systole. The corresponding estimated velocity field is superimposed. In contrast to TDI, that measures mainly the longitudinal contraction towards the apex (Fig. 6.5b), the estimated motion field also captures the significant inward motion of the cardiac walls. This corresponds well to the expert echocardiographic reading.

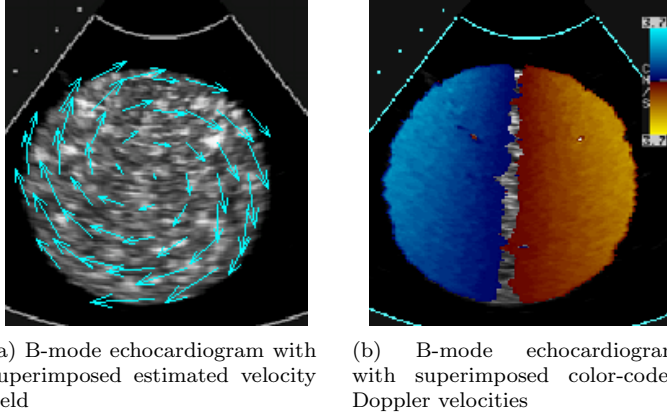




Figure 6.3: Frame of rotating phantom echocardiogram with superimposed velocity information.



(a) Angular error for synthetic ultrasound data

(b) Angular error for phantom sequence

Figure 6.4: Average angular error $\bar{\theta}$ for different values of λ computed for synthetic and phantom ultrasound data.



(a) B-mode echocardiogram with superimposed estimated velocity field

(b) B-mode echocardiogram with superimposed color-coded Doppler velocities

Figure 6.5: Frame of clinical echocardiogram with superimposed velocity information during systole.

6.6 Conclusions

We have proposed a new method to estimate heart motion from echocardiograms that combines information from B-mode and tissue Doppler modalities. Validation experiments on synthetic and phantom data demonstrate that the inclusion of tissue Doppler measurements significantly improves the motion field accuracy. First applications to clinical echocardiograms also give realistic results. Unfortunately, application in clinical routine is still limited since the simultaneous acquisition of B-mode and tissue Doppler with current ultrasound systems is only feasible at a lower spatial and temporal resolution than operating in a single mode. However, the newest generation of ultrasound systems will greatly reduce this limitation, opening the perspective for an extensive validation of the algorithm in clinical practice.

When your heart speaks, take good notes.

— Judith Campbell

Chapter 7

Strain Rate Analysis

Abstract — In addition to ventricular inward motion, systolic myocardial contraction is characterized by a significant wall thickening and circumferential/longitudinal shortening. In this chapter, we present a method to analyze regional myocardial deformation in terms of strain rate. In particular, we compute the magnitudes and principal directions of regional deformation from the estimated parameters of the local affine motion model. For visualization, the principal deformation directions and magnitudes are superimposed onto the echocardiograms in the form of small ellipses. The display provides a means to identify non-contracting regions of the myocardium more easily. The method is tested on synthetic ultrasound data and its potential diagnostic value is demonstrated on clinical echocardiograms.

7.1 Introduction

Local myocardial velocity is an important feature to assess myocardial function. However, it does not allow to differentiate actively contracting tissue from infarcted one that merely moves along with neighboring healthy segments. During normal physiological contraction, the pumping capacity of the ventricle is usually raised by an additional wall-thickening and a longitudinal contraction of the basal segments towards the apex. Due to the wall thickening, the velocity of the endocardium (inner border) is higher than that of the epicardium (outer




Figure 7.1: Deformation of a one-dimensional object of length L . The end points x_0 and x_1 move with velocities $u(x_0, t)$ and $u(x_1, t)$, respectively.

border), but the difference is not necessarily related to the underlying wall motion. Thus, the spatial velocity gradient—also known as *strain rate* [80]—can be of great diagnostic value to differentiate active from passive tissue.

At present, local strain parameters are obtained from ultrasound M-mode recordings and, recently, from tissue Doppler imaging (TDI). However, these methods are essentially one-dimensional as they are restricted to the ultrasonic beam direction. An attractive feature of the proposed local affine motion model (4.2) is that it gives also access to regional two-dimensional strain rate information. The spatial affine parameters of the local motion model form a local strain-rate tensor from which we compute the principal directions and magnitudes of contraction and expansion.

This chapter is organized as follows: in Section 7.2, we review the principles of strain and strain rate as they are used in current ultrasound systems. Then we propose the two-dimensional strain rate analysis method in Section 7.3. Experimental results on synthetic and clinical ultrasound data are presented in Section 7.4.

7.2 Definition of Strain and Strain Rate

Strain defines the amount of deformation of an object caused by an applied force. There are two ways to calculate strain, depending on the reference system used. As sketched for the one-dimensional case in Fig. 7.1, the so-called *Lagrangian* strain $\epsilon_L(t_1)$ at time instance t_1 is defined as the relative elongation with respect to the initial length $L(t_0)$, i.e.,

$$\epsilon_L(t_1) = \frac{L(t_1) - L(t_0)}{L(t_0)}. \quad (7.1)$$

The *natural* or *Eulerian* strain ϵ_N is defined as

$$\epsilon_N(t_1) = \int_{t_0}^{t_1} d\epsilon_N(t), \quad (7.2)$$

where

$$d\epsilon_N(t) = \frac{L(t+dt) - L(t)}{L(t)} \quad (7.3)$$

is an infinitesimally amount of deformation occurring during the infinitesimally time interval dt . In contrast to Lagrangian strain, the reference length is not constant over time but is given by the instantaneous length $L(t)$. By rewriting $d\epsilon_N(t) = \dot{L}(t)/L(t) dt$ it follows that Lagrangian and natural strain are related as $\epsilon_L(t_1) = e^{\epsilon_N(t_1)} - 1$.

The instantaneous rate of deformation—the so-called *natural strain rate*—is defined as

$$\frac{d\epsilon_N(t)}{dt} = \frac{\dot{L}(t)}{L(t)}. \quad (7.4)$$

Using the definition $L(t) = x_1(t) - x_0(t)$ and the fact that $\dot{L}(t) = \dot{x}_1(t) - \dot{x}_0(t) = u(x_1(t), t) - u(x_0(t), t)$, it follows that

$$\frac{d\epsilon_N(t)}{dt} = \frac{u(x_1(t), t) - u(x_0(t), t)}{x_1(t) - x_0(t)}. \quad (7.5)$$

For $\lim_{x_1 \rightarrow x_0}$, this converges to the spatial derivative $\left. \frac{\partial u}{\partial x} \right|_{x_0(t)} = u_x|_{x_0(t)}$ of the instantaneous velocity $u(x, t)$. Thus, strain rate can be interpreted as the speed at which tissue deformation (i.e., strain) occurs; it is measured in (cm/s)/cm = 1/s units. Strain and strain rate relate to each other as displacement does to velocity.

Current approaches to calculate myocardial strain rate are based on tissue Doppler imaging [81, 82], where the axial strain rate component is computed as the spatial derivative of the Doppler velocities as illustrated in Fig. 7.2. Alternatively, one-dimensional strain parameters can also be obtained from ultrasound M-mode recordings [83, 84]. In [84], the local displacement is first estimated from the radio frequency signal by using a cross-correlation method. Strain is then derived as the spatial derivative of the displacement and strain rate is obtained by computing the temporal derivative of the strain data. A review of strain rate principles and existing measurement techniques can be found in [80].




Figure 7.2: Computation of the axial strain rate $(v(r_0 + \Delta r) - v(r_0))/\Delta r$ from tissue Doppler velocities $v(r)$.

Since existing methods are limited to a few selected regions of interest and highly depend on the insonification direction, it is highly desirable to extend the strain rate analysis to two dimensions.

7.3 Two-Dimensional Strain Rate Analysis

In the case of our motion analysis method, two-dimensional strain rate is inherently contained in the underlying local motion model. We recall that the spatial affine model (6.2)

$$\begin{pmatrix} u(x, y) \\ v(x, y) \end{pmatrix} = \begin{pmatrix} u_0 \\ v_0 \end{pmatrix} + \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix} \cdot \begin{pmatrix} x - x_0 \\ y - y_0 \end{pmatrix}$$

is defined within a two-dimensional spatial window at a single time t_0 . The vector $(u_0, v_0)^T$ corresponds to the velocity at the window center $(x_0, y_0)^T$, and the spatial velocity derivatives u_x , u_y , v_x , and v_y are nothing but the strain rate parameters which are assumed to be constant within the local neighborhood.

To exploit the coherence of motion in time, we also added a temporal model component. The spatio-temporal-affine velocity model (4.2)

$$\begin{pmatrix} u(x, y, t) \\ v(x, y, t) \end{pmatrix} = \begin{pmatrix} u_0 \\ v_0 \end{pmatrix} + \begin{pmatrix} u_x & u_y & u_t \\ v_x & v_y & v_t \end{pmatrix} \begin{pmatrix} x - x_0 \\ y - y_0 \\ t - t_0 \end{pmatrix}$$

$$= \begin{pmatrix} u_0 \\ v_0 \end{pmatrix} + \begin{pmatrix} u_t \\ v_t \end{pmatrix} (t - t_0) + \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix} \begin{pmatrix} x - x_0 \\ y - y_0 \end{pmatrix}$$

uses a three-dimensional window function in space-time. It assumes the deformation to be stationary in time but accounts for a temporal change of the translational motion component. In principle, the local motion parameters of this model are more robust since they are estimated within a larger neighborhood containing several frames.

The Jacobian matrix of spatial velocity derivatives,

$$\mathbf{J} = \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix}, \quad (7.6)$$

is also known from mechanical engineering as *strain rate tensor*. It can be decomposed into two terms which are symmetric and antisymmetric, respectively [85]:

$$\mathbf{J} = \mathbf{R} + \mathbf{D}, \quad (7.7)$$

where

$$\mathbf{R} = \frac{1}{2}(\mathbf{J} - \mathbf{J}^T) = \begin{pmatrix} 0 & -\frac{v_x - u_y}{2} \\ \frac{v_x - u_y}{2} & 0 \end{pmatrix} \quad (7.8)$$

and

$$\mathbf{D} = \frac{1}{2}(\mathbf{J} + \mathbf{J}^T) = \begin{pmatrix} u_x & \frac{v_x + u_y}{2} \\ \frac{v_x + u_y}{2} & v_y \end{pmatrix}. \quad (7.9)$$

The matrix \mathbf{R} corresponds to a rigid, rotational velocity field that leaves the local tissue area unchanged. The angular velocity is given by $\omega = (v_x - u_y)/2$ which corresponds to one-half of the curl of the velocity field. The second term, \mathbf{D} , accounts for the deformation of the heart tissue, both contraction/expansion and shear. The components on the principal diagonal describe a dilation along the coordinate axes, whereas the off-diagonal components correspond to a shearing. Since \mathbf{D} is symmetric, it has real eigenvalues λ_1, λ_2 and orthonormal eigenvectors \mathbf{u}_1 and \mathbf{u}_2 ; it can be decomposed as

$$\mathbf{D} = \mathbf{U} \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \mathbf{U}^T, \quad (7.10)$$

where

$$\mathbf{U} = (\mathbf{u}_1, \mathbf{u}_2) = \begin{pmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{pmatrix} \quad (7.11)$$

is the rotation matrix specified by the eigenvectors, and where ϕ denotes the rotation angle of the coordinate system. Thus, the deformation matrix \mathbf{D} can be interpreted as a pure contraction/expansion along the directions of the eigenvectors. Negative eigenvalues λ_i describe contraction, while positive eigenvalues correspond to dilation of the heart tissue. The divergence, given by the trace of \mathbf{D} (or equivalently by the sum of its eigenvalues), describes the local area change of the tissue.

7.3.1 Incompressibility Assumption

Like most soft tissues, the myocardium is assumed to be nearly incompressible [86, 50]. Volume changes due to blood flow are considered to be negligible. Although this assumption is only true in three dimensions, it may be applicable in two dimensions if the deformation orthogonal to the image plane is small. In particular, this condition is satisfied when the ventricle is imaged orthogonal to the fiber directions because the tissue elasticity along fiber directions is much lower than across fibers. As known from continuum theory, the divergence of a velocity field of an incompressible medium must be zero [85]. In the two-dimensional case, this reads

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0. \quad (7.12)$$

Since we use a local affine model for the velocity, whose spatial derivatives u_x , u_y , v_x , and v_y are assumed to be constant, the divergence-free constraint reduces to

$$u_x = -v_y. \quad (7.13)$$

This simply means that one of the model parameters, u_x or v_y , can be eliminated from the local linear systems (4.5) or (6.5). A direct consequence of the incompressibility assumption is also that the sum of eigenvalues of the deformation matrix \mathbf{D} is zero, i.e., $\lambda_1 = -\lambda_2$.

Besides the fact that this adapted model matches the underlying biomechanical model more closely, it will also increase the robustness to noise of the algorithm because there are fewer parameters to estimate.

7.3.2 Strain Rate Visualization

For visualization, the principal deformation directions and magnitudes are superimposed in the form of small ellipses onto the echocardiograms inside the time varying ROI as shown in Fig. 7.3. The semi-axes directions correspond to




Figure 7.3: Strain rate visualization by means of ellipses. The semi-axes of a circle with radius r_0 are scaled according to the corresponding eigenvalues of the deformation matrix \mathbf{D} .

the eigenvectors of \mathbf{D} and their lengths are given by

$$\left(1 + \frac{\lambda_i}{\lambda_{\text{Max}}}\right) r_0, \quad i = 1, 2,$$

respectively. The normalization parameter λ_{Max} is the maximum absolute value of all eigenvalues computed within the ROI. The parameter r_0 corresponds to the radius of the non-deformed circle and controls the overall size of the ellipses. Negative eigenvalues (tissue contraction) lead to an axis-shortening, whereas positive eigenvalues (tissue dilation) correspond to an elongation of the corresponding semi-axis.

7.4 Numerical Results

7.4.1 Application to Synthetic Data

To test the feasibility of the estimation of strain rate information from image sequences, we first applied our algorithm to the synthetic ultrasound data introduced in Section 4.4.1. These sequences simulate a uniformly contracting and expanding annulus that also undergoes a global rigid translation in the diagonal direction. The applied wall thickening and thinning was chosen such that it satisfies the incompressibility constraint. The maximum strain rate applied was




Figure 7.4: Estimated velocities and strain rates from synthetic ultrasound data during diastole (a), (b) and systole (c), (d). Strain rate is independent from underlying rigid translation. (Note that the ellipses are plotted on a fixed, uniform grid, leading to an asymmetric appearance of the radially symmetric strain.)

0.025 per frame. If a frame rate of 50 Hz is assumed, this value corresponds to the strain rate 1.25 s^{-1} , which is in the range of typical values measured in normal patients [87]. The estimated motion fields and strain rates for one sequence are shown in Fig. 7.4. Fig. 7.4a and 7.4c display the velocity field during diastole and systole, respectively. Since the motion field captures the superposition of translational motion, radial outward/inward motion and deformation, its interpretation remains difficult; for instance, one may get the wrong impression that the heart model contracts less in the lower left part. In contrast, the estimated

strain rate is independent from the underlying rigid translation and represents well the uniform deformation of the annulus. The myocardial thinning during expansion is clearly indicated by the circumferentially elongated ellipses in Fig. 7.4b; likewise, the myocardial thickening during contraction is represented by the radially dilated ellipses in Fig. 7.4d. The fact that the applied deformation close to the inner boundary is larger than at the outer boundary is also indicated by the different ratios of long to short semi-axis lengths.

The above results were computed using the spatio-temporal-affine motion model without incompressibility assumption. To test the accuracy and robustness of the local motion models, i.e., spatio-temporal-affine and spatial affine, and their divergence-free adaptations, we applied them to echo sequences that were obtained from point scatterer images of signal-to-noise ratios varying from 26.02 to 13.98 dB. We used the same multiresolution strategy in all cases to be comparable. To quantify the performance, we use two different error measures. The relative amplitude error between the correct eigenvalues λ_1 and their estimates $\hat{\lambda}_1$, which is computed as

$$E_\lambda = \frac{\|\hat{\lambda}_1 - \lambda_1\|_{l_2}}{\|\lambda_1\|_{l_2}}, \quad (7.14)$$

where the norm is computed over all pixels in the sequence with non-zero motion. The second is the angular root-mean-square error of the principal direction of deformation calculated as

$$E_\phi = \frac{1}{\sqrt{N}} \|\hat{\phi} - \phi\|_{l_2}, \quad (7.15)$$

where ϕ and $\hat{\phi}$ denote the correct and estimated angle, respectively. The norm is again computed over all N pixels with non-zero motion.

The experimental results obtained are plotted in Fig. 7.5. Fig. 7.5a shows the relative amplitude error E_λ of the different methods. The spatio-temporal-affine method is clearly more robust against noise than the spatial affine version since it is based on multiple frames in time. Consequently, the performance improvement of the spatio-temporal-affine method due to the additional incompressibility constraint is much less significant than for the spatial affine one. The average angular error of the principal deformation direction in Fig. 7.5b shows a similar behavior. In this experiment, the performance gain due to the additional divergence-free constraint is relatively low for both motion models.

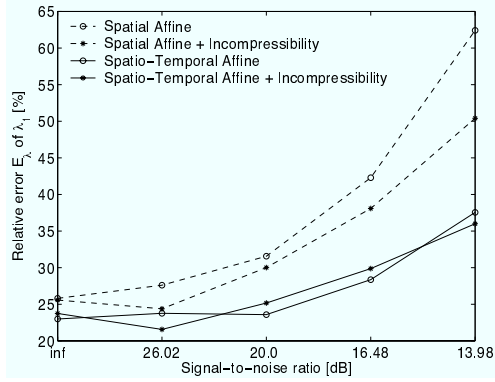

(a) Relative amplitude error E_{λ} of λ_1 (b) Angular root-mean-square error E_{ϕ} of principal direction

Figure 7.5: Average amplitude and directional errors of the estimated strain rate parameters for different motion models. The errors are plotted against different signal-to-noise ratios of the synthetic ultrasound sequences.




Figure 7.6: Estimated strain rates from ultrasound data acquired during an animal study. The hyperactive wall thinning during diastole (a) and thickening during systole (b) in the basal anteroseptal segment (lower right) is clearly indicated.

7.4.2 Application to Clinical Data

The applicability of two-dimensional strain rate analysis to real ultrasound was first tested on clinical echocardiograms obtained from an animal study as used in Section 4.4.3. Here, we used the spatio-temporal-affine motion model without incompressibility constraint. Fig. 7.6 illustrates the estimated strain rates of one case after an artificially induced infarction in the apical anterior septal segment (upper right). Figures 7.6a and 7.6b correspond to one frame during diastole and systole, respectively. According to the expert echocardiographic reading, the mid to basal anterior septal segments (middle right) exhibit a hyperactive wall thinning and thickening to compensate for the apical dyskinesia. This behavior is well indicated by the ellipses which are elongated in the corresponding directions.

The ability to estimate myocardial deformation in clinical echocardiograms acquired during routine clinical examinations is demonstrated in the short axis view shown in Figure 7.7. The circumferential alignment of the ellipses in Fig. 7.7a documents the typical wall thinning and circumferential lengthening during diastole. In contrast, the ellipses are elongated and aligned radially in Fig. 7.7b, indicating myocardial wall thickening and circumferential shortening during systole.




Figure 7.7: Estimated strain rates from apical short axis view. Myocardial radial wall thinning and circumferential lengthening during diastole is indicated by the deformation map (a). The reversed deformation during systole is shown in (b).

7.5 Conclusions

The experiments demonstrate that the proposed strain rate analysis method has the potential to measure myocardial deformation and to identify regions of abnormal contractility. In contrast to existing one-dimensional strain rate imaging, which operates along pre-defined directions, the proposed algorithm is able to determine both the magnitudes and principal directions of the deformation automatically. First applications on clinical echocardiograms give promising results. However, cardiac deformations can in general not be modeled as a two-dimensional system. Consequently, the method will not be accurate for general three-dimensional cardiac deformations that are imaged in a cross-sectional manner. Although the proposed method can be extended to three dimensions in a straightforward way, its application to real echocardiograms is limited since current ultrasound systems are not yet able to provide volumetric data at a sufficiently high frame rate.

*Nobody has ever measured, even poets, how much
the heart can hold.*

— Zelda Fitzgerald

Chapter 8

Clinical Validation

Abstract — In this chapter, we validate the proposed motion analysis algorithm with tissue Doppler and demonstrate its diagnostic potential on a large set of clinical echocardiograms. Correlation between measured myocardial inward motion and conventional subjective contractility assessment is shown by analysis of variance and a subsequent post hoc multiple comparison test. The ability to automatically derive diagnoses based on measured motion data is demonstrated using an artificial neural network classifier. Clinical usefulness of the proposed algorithm is also highlighted on some exemplary case studies.

8.1 Introduction

In the validation examples presented so far, the proposed motion analysis approach—which we also refer to as *Multiscale Motion Mapping* in the clinical context—was tested in setups for which true motion was known; i.e., synthetic ultrasound data and moving physical phantoms. A first clinical exploration was done on a series of prototypical echocardiograms and in an animal study. However, quantitative motion assessment is only an intermediate step towards the ultimate goal of obtaining more objective and reproducible diagnoses. To assess the diagnostic value of the proposed algorithm, validation on a wide range of clinical data is indispensable. Therefore, we evaluated the feasibility of the de-

veloped method in routine echo on a consecutive series of 114 echocardiograms acquired for clinical indications. Conventional reading was done by three experts who assigned subjective ratings ranging from “hyper-” to “dyskinetic” to individual myocardial segments and gave each echocardiogram an overall label as “normal” or “abnormal”. Objective analysis by Multiscale Motion Mapping was performed by two observers who were blind to the subjective analysis.

In a first step, we validated velocity measurements by Multiscale Motion Mapping with tissue Doppler velocities on a subset of echocardiograms for which B-mode and Doppler data were acquired simultaneously. Linear regression and Bland-Altman analysis [88] were used to demonstrate agreement between both methods.

The diagnostic value of the measured velocity data was validated in a second step. First, we evaluated the correlation between the extent of wall motion abnormalities, as rated by the experts, and radial velocity magnitudes in individual segments. Correlation between measured inward motion and subjective rating was validated by one-way analysis of variance (ANOVA) and a subsequent post hoc multiple comparison test. Then, velocity measurements in different myocardial segments and their corresponding subjective ratings were combined and used to train an artificial neural network by using the first half of the available data. The ability to automatically classify patients globally as normal or abnormal based on the measured velocities was evaluated using the second half of the data.

8.2 Experimental Setup and Statistical Validation

We collected prospectively a consecutive series of clinical echocardiograms from 125 patients (mean age 64 ± 14 years) referred for echocardiography for clinical investigation. Since a complete set of standardized two-dimensional echocardiographic views (cf. Section 3.3 on page 27) was not available in all patients, we analyzed the two- and four-chamber long axis views in our experiments. These two orthogonal views are most frequently acquired in clinical practice. Image quality was deemed sufficient for quantification in 114 patients (91%). Thereof, 38 patients (33%) were rated as normal by the experts.

Multiscale motion analysis was applied independently by two blinded users and produced completely reproducible motion maps. For the multi-scale refinement scheme, described in Section 4.2.3, we used a five-level pyramid decompo-

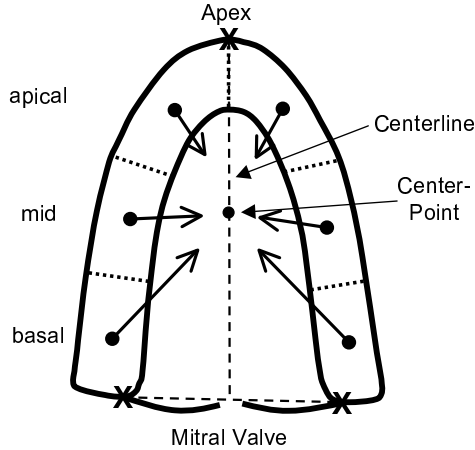


Figure 8.1: Radial velocities with respect to a center-point are extracted at six locations defined by a clinical observer. The center-point corresponds to the midpoint of the longitudinal centerline defined by means of three user-defined anatomical landmarks (apex and two mitral valve attachments).

sition ($k = 0, \dots, 4$). To extract clinically relevant motion information, the two observers specified subjectively six locations in each of the two views analyzed. The locations were chosen corresponding to the standardized myocardial segmentation introduced in Section 3.3.3 on page 29. A schematic representation is shown in Fig. 8.1. Velocities were extracted for the mid-systole frame that was determined by means of the ECG-signal. By analogy with Section 4.3, the velocity vectors at the selected positions were projected onto the radial direction defined with respect to the center-point of the ventricle cross-section. The center-point was defined as the midpoint of a longitudinal centerline connecting base and apex. The centerline was determined by means of three anatomical landmarks, the two basal mitral valve attachments and the apex. We chose to analyze the radial velocity component because it captures the typical inward/outward motion pattern of the ventricle; the use of the magnitude of the two-dimensional velocity vectors is less appropriate because it does not allow to distinguish between the motion orientation. Other velocity components such as motion parallel or perpendicular to the centerline also provide potential diagnostic information but were not analyzed in this study.

Observer dependency of velocity extraction was analyzed by computing the

mean and standard deviation of the difference between the two sets of measurements. The interobserver variability was 0.896 cm/s (5.1% of measured range), with a mean velocity difference of -0.077 cm/s indicating a non-significant bias ($p = 0.04$). The correlation coefficient between the two observers was 0.83 ($p < 0.001$). The intraobserver variability of repeated velocity extraction on a subset of 28 echocardiograms by the first observer was 0.568 cm/s (3.2% of measured range) with a mean velocity difference of -0.063 cm/s ($p=NS$). The corresponding correlation coefficient was 0.93 ($p < 0.001$). Note that this variability is entirely due to the manual placement of landmarks and not to the global motion field which is fully reproducible. The measurements show a good agreement and reproducibility of the manual velocity extraction. A completely automatic extraction could entirely remove the observer dependency and increase the consistency of subsequent classification algorithms but is not considered here.

In parallel, three experienced echocardiographers who were blind to the quantitative motion analysis rated individual segments in each echocardiographic loop for visible motion in a semi-quantitative manner. We used a scale from 3 to -1 , indicating dyskinesia (3), akinesia (2), hypokinesia (1), normokinesia (0), or hyperkinesia (-1) as defined in Section 2.4 on page 17. In addition, the experts assigned the overall, binary label normal or abnormal to each patient.

8.2.1 Comparison with Tissue Doppler

Agreement between Multiscale Motion Mapping and tissue Doppler velocities was assessed in a subset of 11 echocardiograms for which B-mode and tissue Doppler signals were acquired simultaneously. In each loop, six segmental velocities were extracted from one systolic and one diastolic frame, respectively. For comparison, only the motion vector component parallel to the ultrasound beam directions was used because tissue Doppler is unable to measure the motion component perpendicular to the beam. As shown in Fig. 8.2a, there was a good correlation between Multiscale Motion Mapping and Doppler, with a correlation coefficient of 0.94 (95% confidence interval: [0.92, 0.96]).

Bland-Altman analysis [88] of the method comparison is illustrated in Fig. 8.2b. The Bland-Altman graph consists of plotting the pairwise velocity means versus their differences. The mean of the velocity differences was 0.223 cm/s (95% confidence interval: [-0.039, 0.486]), corresponding to no statistically significant bias. The standard deviation of the velocity differences was 1.527 cm/s (7.2% of measured range).

Table 8.1: Critical p -values of F -test for the null hypothesis that all velocity samples in a given myocardial segment are drawn from the same population.


Longitudinal position	Two-chamber view		Four-chamber view	
	Inferior	Anterior	Septal	Lateral
Apical	0.017	0.224	0.007	0.009
Mid	< 0.001	0.017	< 0.001	< 0.001
Basal	0.002	0.166	< 0.001	< 0.001

Both graphs document good agreement between the two approaches. They indicate that deviation of Multiscale Motion Mapping from tissue Doppler grows with increasing velocity magnitudes. This is most probably due the the lack of a sufficiently high frame rate in current ultrasound systems. The problem will diminish with the advent of new ultrasound systems that provide frame rates between 75 to 100 Hz compared to about 25 to 50 Hz of current systems.


8.2.2 Segmental Multiple Comparison Tests

For clinical validation, we first analyzed the twelve myocardial segments individually to investigate whether the measured velocities correlate with the subjective expert rating. For each segment, we performed a one-way analysis of variance (ANOVA) [89] to compare the velocity means of the different classes; i.e., hyper-, normo-, hypo-, a-, and dyskinetic. We used the Fisher (F) statistic to test the null hypothesis that there are no differences between the velocity means for different classes, suggesting that the variance of the within-class samples should be identical to that of the between-class samples. If $F \gg 0$, it is likely that differences between class means exist. The critical p -values for the individual segments are summarized in Table 8.1. In this study, a p -value of 0.05 was regarded as significant. The results show that velocity means of different disease classes are significantly different in ten segments. The test was non-significant in the apical and basal anterior segments (two-chamber view).

The Fisher statistic alone, however, cannot indicate which of the means is responsible for rejecting the null hypothesis that class means are equal. To investigate our results in more detail, we applied a post-hoc multiple comparison test; in particular, we used the Scheffé test [90] which tests all possible pairs of classes for differences between means. The Scheffé test is known to be



(a) Linear regression (Fitted line: $y = 0.87x - 0.14$)



(b) Bland-Altman plot (Mean: 0.223 cm/s, SD: 1.527 cm/s)

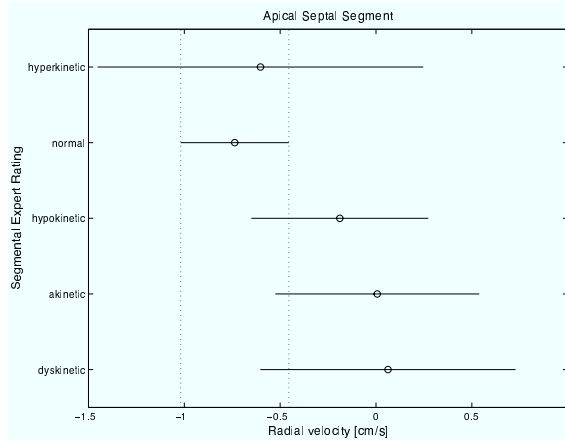
Figure 8.2: Regression (a) and Bland-Altman analysis (b) of tissue Doppler and Multiscale Motion Mapping-based velocity measurements.

conservative, which helps to reduce spurious detections.


The Scheffé test results for the basal, mid and apical segments of the septal wall (four-chamber view) are shown in Fig. 8.3 for illustrative purposes. Subjective classifications were taken from the first expert in this case. Fig. 8.3a shows the velocity means of the different classes along with so-called uncertainty intervals [90, p. 96 et sqq.]. Two velocity means can be declared significantly different from each other if and only if their uncertainty intervals do not overlap. Negative and positive velocities represent inward and outward motion, respectively. The graph indicates that significant differences between velocity means could not be measured in the apical segment. In this case, only a slight trend is visible. This is probably due to the fact that the apical part of the myocardium moves very little, which makes visual assessment and classification into multiple categories difficult. In addition, image quality in the near field of the ultrasound beam is usually worse than in the far field. This results into a relatively high noise level which degrades the accuracy of velocity measurements in this area. However, the situation is much different in the mid and basal segments. As shown in Fig. 8.3b, velocity means of normal, hypokinetic and akinetic mid-segments are mutually significantly different in this region. Here, no segments were classified as hyper- or dyskinetic. In accordance with visual reading, mean radial velocities of normal mid septal segments are higher than in the apical part. The highest velocities are observed in the basal segments (Fig. 8.3c). The velocity mean of normal basal segments exhibits a significant difference from abnormal ones. A sub-division of abnormal motion into hypo- or akinesia is not significant here.

Post hoc tests in the inferior, anterior and lateral ventricular walls showed similar results. In summary, velocity means between normal and abnormal segments were significantly different in all mid and basal segments. Only in the basal anterior region (two-chamber view), normal segments could not be distinguished from abnormal ones ($p = 0.166$). Velocity means of different sub-classes of abnormal motion (hypokinetic, akinetic, dyskinetic) were only significant in a few cases. Significant velocity differences between normal and abnormal motion in the apical regions could not be established; nevertheless, a trend was always visible.

Altogether, the test results suggest that radial velocity could serve as a useful feature to discriminate and classify myocardial segments based on the velocity measurements. The performance of a neural network-based classifier is investigated in the following section.




(a) Apical septal



(b) Mid septal

Figure 8.3: Post hoc Scheffé test: Means of radial velocities grouped according to visual expert rating are plotted along with their corresponding uncertainty intervals. Non-overlapping intervals indicate that differences between group means are statistically significant. Negative and positive velocities represent inward and outward motion, respectively.



(c) Basal septal

Figure 8.3: Post hoc Scheffé test: Means of radial velocities grouped according to visual expert rating are plotted along with their corresponding uncertainty intervals. Non-overlapping intervals indicate that differences between group means are statistically significant. Negative and positive velocities represent inward and outward motion, respectively.




Figure 8.4: Neural network setup for automatic patient classification. The twelve extracted radial velocity measurements are fed into the input layer. The output layer corresponds to subjective segmental motion ratings and an overall classification of a patient as normal or abnormal.

8.2.3 Artificial Neural Network Classification

The extracted velocity measurements and corresponding subjective ratings were used to train an artificial neural network [91], as indicated in Fig. 8.4. The twelve radial velocity measures for each patient (six from the two- and four-chamber view, respectively) were used as input data. The output layer corresponded to the subjective expert ratings of segmental motion and the overall label classifying a patient as “normal” or “abnormal”. Two networks were trained separately for each observer who extracted radial velocity measurements to study the influence of manual velocity extraction on the classification. In both cases, the available data was split into two sets: the first half of the patients was used to train the network while the second half was used for testing. For training, we used the majority vote of the three experts in the output layer. For the network design, we used a three-layer structure, sigmoid transfer functions and the error-backpropagation algorithm for learning. Network outputs were rounded to the nearest integer to determine the corresponding class.

The quality and usefulness of the neural network-based classification was assessed using several characteristics that are commonly used in medical diagnostic tests; in particular, we used the following parameters:

Sensitivity: The sensitivity $TP/(TP + FN) \in [0, 1]$ of a test is the probability

that it will produce a true positive (TP) result on a pathological population as compared to a reference or “gold standard”. FN denotes false negative ratings.

Specificity: The specificity $TN/(TN + FP) \in [0, 1]$ is the probability that a test will produce a true negative (TN) result when applied to normal (non-pathological) patients. FP denotes false positive ratings.

Positive Predictive Value: The positive predictive value $TP/(TP + FP) \in [0, 1]$ is the probability that a person is pathological when a positive test result is observed.

Negative Predictive Value: The negative predictive value $TN/(TN + FN) \in [0, 1]$ is the probability that a person is normal when a negative test result is observed.

Diagnostic accuracy: The diagnostic accuracy $(TP+TN)/(TP+TN+FP+FN) \in [0, 1]$ is the probability that a test will produce a true result. It is an important measure of the overall performance of a test.

Kappa Coefficient: The kappa coefficient $-1 \leq \kappa = (P_o - P_e)/(1 - P_e) \leq 1$ indicates the mutual agreement among raters who make category judgments [92]. It expresses the ratio between the observed concordance frequency P_o and the maximum possible agreement. Both frequencies are corrected by the frequency P_e denoting agreement by chance. The value $\kappa = -1$ corresponds to complete disagreement, $\kappa = 0$ indicates that agreement is just due to chance, and $\kappa = 1$ denotes perfect agreement between ratings. More specifically, the guidelines for interpreting kappa are as follows [93]: $\kappa \leq 0.4$: poor, $0.4 < \kappa \leq 0.6$: fair, $0.6 < \kappa \leq 0.74$: good, and $0.74 < \kappa \leq 1$: excellent agreement.

In order to assess the performance of the semi-automatic classification, we evaluated its ability to classify patients globally as normal or abnormal. In our experiments, the diagnostic accuracy between objective classification and the majority vote of the experts was 84% for both observers who extracted velocity measurements. In both cases, the sensitivity was larger than the specificity. The results for all above-mentioned test characteristics are summarized in Table 8.2.

For conventional, subjective classification, kappa values for interobserver agreement between two different experts were 0.55, 0.59 and 0.71, respectively. This corresponds to a fair-to-good agreement. The multi-rater kappa coefficient

Table 8.2: Test statistics for semi-automatic classification vs. majority consensus of expert ratings that a patient is globally normal or abnormal.

Test Characteristic	User 1 vs. Experts	User 2 vs. Experts
Sensitivity	0.91	0.88
Specificity	0.60	0.70
Positive predictive value	0.88	0.91
Negative predictive value	0.67	0.64
Diagnostic accuracy	0.84	0.84
Kappa coefficient	0.53	0.56

of concordance between all three experts was 0.62. The kappa coefficients between neural network-based classification and subjective majority rating by the experts were 0.53 and 0.56, respectively. This is in the same range as the level of mutual expert agreement. The interobserver agreement for the two sets of semi-automatic classifications was good-to-excellent ($\kappa = 0.74$), meaning that the semi-automatic algorithm is more reproducible than pure subjective rating.

The validation results demonstrate that the proposed quantitative motion analysis has in fact some potential to provide more objective and reproducible diagnoses. In this study, classification was only derived from a few radial velocity measurements in a single frame during systole. The diagnostic accuracy could probably be further increased by extending the analysis over several frames in time and by also including more cross-sectional views than just two. Grouping together myocardial segments that are perfused by the same coronary artery (LAD, LCX or RCA, cf. Section 2.2.4 on page 14) could also increase classification consistency. In addition, strain rate measurements provide complementary information about cardiac contractility which may improve subsequent classification algorithms significantly.

8.3 Exemplary Case Studies

After statistical validation of the proposed quantitative motion assessment on a large set of echocardiograms, we illustrate its clinical value on a few exemplary cases to highlight its potential and challenges.

8.3.1 Comprehensive Objective Analysis of Paradoxical Motion

To exemplify the practical application of Multiscale Motion Mapping, an example of a routine echocardiogram, for which local echo experts disagreed as to the presence or absence of motion abnormalities, is given in Fig. 8.5. Fig. 8.5a and Fig. 8.5b show the native gray scale loop together with selected motion vectors in the septum and the lateral wall during diastole and systole, respectively. They indicate clearly that there is paradoxical motion of the basal interventricular septum. Additional insight can be gained through separate analysis of transverse (towards the centerline) and longitudinal (along the centerline) ventricular motion. The motion components at mid-systole indicate that there is reduced longitudinal velocity of the basal segments ($p < 0.01$ vs. normal). Transverse velocity of the lateral segments was in the normal range, whereas the basal septum moves outwards. As indicated in Fig. 8.5c and Fig. 8.5d, two-dimensional strain rates show normal behavior during diastole and systole, respectively. Measured peak systolic thickening/thinning strain rates in the basal septum were 1.28 s^{-1} and -1.03 s^{-1} with maximum thickening directed towards the mid-ventricle.

These findings of objective, quantitative analysis—namely, paradoxical motion but normal thickening of the septum—agree nicely with careful subjective analysis of the echocardiogram. Such comprehensive, objective evaluation of heart motion is not possible with existing echocardiographic techniques such as Doppler.

8.3.2 New Insights into Cardiac Motion

Rotational motion of the left ventricular apex is a phenomenon that is frequently observed visually by physicians in echocardiograms. The rotational or torsional motion has been poorly understood and has lacked clinical relevance. It has been postulated that apical twisting characterizes normal systolic function in humans, but it has been difficult to measure it noninvasively. Studies of animal models have demonstrated that a decrease in the amplitude of apex rotation correlates strongly with the presence of ischemia [94].

The ability of Multiscale Motion Mapping to assess apical myocardial twisting is illustrated in Fig. 8.6. The motion vectors in Fig. 8.6a document the clockwise apical twisting during early diastole. The complexity of the actual motion pattern becomes readily apparent in the fish-swarm like appearance of the arrows, illustrating the difficulties of motion assessment with conventional




Figure 8.5: Paradoxical motion in an apical four-chamber view. The interventricular septal wall moves inwards during diastole (a) and partially outwards during systole (b). Myocardial diastolic thinning (c) and systolic thickening (d) remain normal.

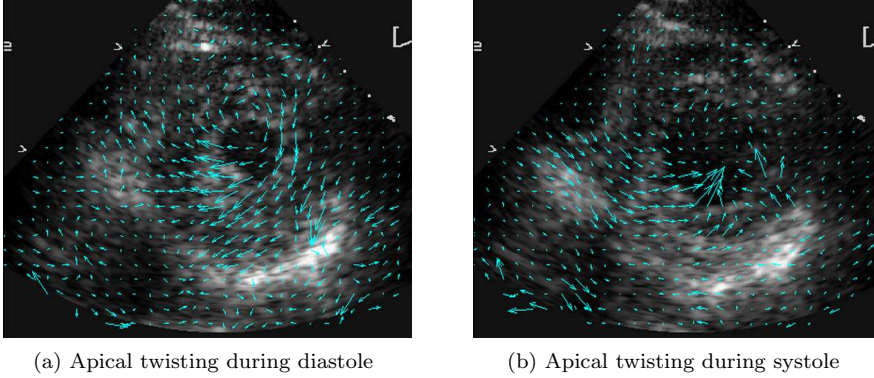


Figure 8.6: Highlighting of myocardial twisting motion in an apical short axis view. Clockwise apical twisting during early diastole is documented by the motion arrows in (a). The reverse twisting motion during systole is indicated in (b).

Doppler-based methods that are limited to detecting a one-dimensional motion component only. The counter-clockwise rotation of the ventricular apex during systole is captured by the motion field shown in Fig. 8.6b. Since the rotational motion component with respect to the ventricular center can be easily extracted from the motion field, further quantitative studies may reveal its diagnostic value to detect and quantify myocardial ischemia or infarction.

8.4 Conclusions

The diagnostic potential of the proposed motion analysis algorithm was investigated. Clinical applicability was tested on a consecutive series of echocardiograms, where good feasibility, fair correlation with expert rating, and good intra- and interobserver concordance were documented. Exploratory case studies documented the usefulness of Multiscale Motion Mapping in a set of complex motion analyses including abnormal septal motion and analysis of myocardial twisting. Automated classification could probably be further improved by using multiple frames, a higher number of cross-sectional views, by taking into account segmental correlation due to coronary perfusion, and by including strain rate measurements into the algorithm. The study also revealed the high intra- and interobserver variability of conventional motion assessment—calling even

stronger for more objective and reproducible approaches.

The only lasting beauty is the beauty of the heart.

— Jelaluddin Rumi

Chapter 9

Conclusion

We presented a novel approach to a quantitative and objective assessment of cardiac motion and deformation from ultrasound. It provides a step towards a more consistent evaluation of cardiac function by overcoming the subjectivity of conventional echocardiographic reading. The main contributions of this thesis are summarized in the following.

9.1 Main Contributions

- We proposed a sliding-window-type algorithm that estimates myocardial motion from B-mode echocardiograms. The applied affine motion model is particularly well suited for describing typical ventricular motions such as contraction/expansion and shear. To increase the robustness to noise, we added a local linear model for the velocity in time, which allows us to base the parameter estimates on a larger sample size. From the affine model parameters, we compute local myocardial deformation in terms of strain rate. In contrast to conventional Doppler-based strain rate imaging, which is restricted to the scan lines, we are able to compute both the principal directions as well as the amounts of deformation.
- To be able to estimate large motions and to further increase the robustness to noise, we apply a multi-scale strategy in space. The motion vectors are

computed in a coarse-to-fine fashion from image pyramids that represent image copies at multiple resolutions.

- We introduced the concept of multiresolution moment filters, a multi-channel filtering algorithm to compute local weighted geometric moments efficiently within windows of dyadic sizes. We demonstrated that B-splines are well suited window functions; besides their nearly isotropic Gaussian shape, they satisfy a two-scale equation—a key feature for efficient multiresolution processing. The algorithm enables us to choose the most appropriate window size adaptively to the underlying motion at a low computational cost. Besides their application in motion analysis, we demonstrated their versatility and usefulness by applying them to two other image analysis problems: image denoising and feature extraction. In these cases, the power of sliding-window-type algorithms in combination with the principles of *weighting* and *multiresolution* became evident.
- Since Doppler-based motion measurements provide valuable, partial motion information, we extended our motion estimation algorithm to exploit data from two modalities: B-mode and tissue Doppler. The incorporation of Doppler further increased the accuracy and robustness of the approach.
- We showed that we can extract clinically relevant information from the estimated motion data and proposed an intuitive display of the results to assist the physician in finding a more objective diagnosis. We chose to display the radial motion component with respect to the ventricular center in color-coded form to directly visualize myocardial inward and outward motion. This display is more straightforward than conventional color Doppler imaging, which displays the velocity components towards the transducer. In addition, we proposed to superimpose regional myocardial deformation information in the form of deforming ellipses that indicate the principal directions and the amount of deformation in a dynamic fashion. For both displays, we developed a robust algorithm to track a user-defined region of interest in time.
- To enable the routine application of the algorithm in clinical practice, we developed a user-friendly graphical user interface. The on-site installation of the software facilitates the execution of further clinical validation studies.
- The proposed algorithm was tested in several fashions. We verified the correctness of the algorithm on synthetic ultrasound data. We demonstrated

the ability of our algorithm to extract motion from real ultrasound data by performing phantom experiments and exemplary clinical case studies. We also documented the diagnostic value of the estimated motion data by a large-scale clinical validation study. The largest study on 114 patients required the order of eight months of work to incorporate the software into the Basel clinical system, to organize and to handle the data, to analyze the data, and finally to perform statistical comparisons between the extracted motion and the conventional motion assessment by several experts, which had also to be coordinated. The measured velocities showed good correlation with conventional diagnostic reading. A first neural network-based classification experiment revealed the potential to derive objective and reproducible diagnoses from the motion data in an automated fashion.

9.2 Future Research Directions

Although the proposed algorithm yields promising experimental results, there is potential to further improve its accuracy and applicability. Some possible research directions are listed below.

- To account for gray value changes due to out-of-plane motion and signal dropouts, the optical-flow constraint may be modified accordingly. For instance, the right hand side could be replaced by a non-zero term that captures brightness changes that are not due to motion. However, a physical model is hard to determine because cardiac motion is very complex and varies strongly in different parts of the myocardium.
- The validity of the local motion model and the brightness conservation along the motion trajectories depend heavily on the frame rate of the ultrasound systems. An increase in acquisition speed—which has already remarkably improved in the newest generation of machines—would further raise the feasibility of the proposed approach.
- As mentioned in Chapter 6, the bimodal acquisition of B-mode and tissue Doppler is still hampered by a significant loss in image resolution. Therefore, the performance of the proposed bimodal motion analysis approach is still limited compared to the analysis of B-mode data alone. An improvement of the bimodal acquisition technique may considerably increase the impact of the valuable Doppler measurements to obtain more accurate and robust motion estimates.

- Instead of analyzing the scan converted and post-processed B-mode data, one may opt to directly analyze the radio-frequency data. The optical-flow-based cost function may be replaced by specifically tailored motion constraints formulated on the RF-data. If the raw data and the various post-processing steps were known, the least-squares estimator used in this work could be adapted/replaced to take the underlying noise model into account. However, all this requires a close collaboration with hardware vendors to have access to the raw image data that is usually only available in a proprietary format.
- Clearly, the first clinical validation needs to be followed by more thorough clinical studies in a wide variety of clinical problems. Clinical trials in the field of routine echo, heart failure quantification and stress echo are presently taking place using our software in Basel.
- The advent of real-time, three-dimensional ultrasound systems opens the door for a more comprehensive, 3D-assessment of motion. Although the proposed motion analysis algorithm is only presented here in two dimensions, it can readily be extended to three dimensions. A volumetric analysis would capture complex heart motions much better and remove problems such as out-of-plane motion in the current two-dimensional assessment. The proposed divergence-free constraint of the local velocity model, which takes into account the physical incompressibility of the myocardium, will also be more consistent in this case.

The three-dimensional extension of the proposed sliding-window-type algorithm would still correspond to solving small linear systems, which can be done efficiently using direct methods and bears a huge potential for parallel implementation. This *divide and conquer* principle is a great advantage over global motion analysis approaches, such as the Horn-Schunk method, that need to solve huge sparse linear systems whose size grows dramatically with increasing dimensions. This usually requires sophisticated iterative methods and an efficient memory handling.

- A fully automated method for analyzing echocardiograms is challenging, but may be achieved by an integration of the proposed method with adequately tuned active contour and trainable shape models. The resulting algorithm could be a weighted combination of the different approaches where the relative weights are based on appropriate confidence measures.

*Nothing is less in our power than the heart, and
far from commanding we are forced to obey it.*

— Jean J. Rousseau

Appendix A

Multiresolution Moment Filters

Abstract — We introduce local weighted geometric moments that are computed from an image within a sliding window at multiple scales. When the window function satisfies a two-scale relation, we prove that lower order moments can be computed efficiently at dyadic scales by using a multiresolution wavelet-like algorithm. We show that B-splines are well suited window functions because, in addition to being refinable, they are positive, symmetric, separable, and very nearly isotropic (Gaussian shape). We present three applications of these multi-scale local moments. The first is a feature extraction method for detecting and characterizing elongated structures in images. The second is a noise reduction method which can be viewed as a multi-scale extension of Savitzky-Golay filtering. The third is a multi-scale optical-flow algorithm that uses a local affine model for the motion field, extending the Lucas-Kanade optical-flow method. The results obtained in all cases are promising.

This chapter is based on our paper [95].

A.1 Introduction

Global geometric moments and their invariants are widely used in many areas of image analysis, including pattern recognition [96], image reconstruction [97], and shape identification [98]. In addition to geometric moments, which are also known as regular or ordinary moments, a number of other moments has been proposed. The notion of complex moments was introduced in [99] for deriving moment invariants. Teague [100] suggested the use of orthogonal moments and introduced complex valued Zernike moments that are defined on a unit disk. A second class of orthogonal moments is given by Legendre moments which make use of Legendre polynomials. The usefulness of Legendre and Zernike moments has been demonstrated, in particular, for image reconstruction [97, 101] and pattern classification [102]. The pseudo-Zernike formulation proposed in [103] further improved these characteristics. A detailed discussion of moment-based image analysis can be found in the monograph [104].

Some authors have applied geometric moments in a local fashion for image and texture segmentation [105], [106] and direction-based interpolation [107]. The idea there was to compute moments locally over some square region of interest which is moved over the image; the window functions may be overlapping or not, depending on the application. An efficient method to compute local moments inside sliding, squared windows with constant weights was recently proposed in [108].

In this paper, we are extending the notion of local geometric moments by introducing two refinements: *weighting* and *multiresolution*. The idea of weighting is motivated by the observation that the square window that has been used so far is rather anisotropic. Indeed, if the goal is to design a “rotation-invariant” algorithm, it makes good sense to apply an isotropic window with a radial weighting that decreases away from the center. Multiresolution is a feature that is highly desirable for designing image processing algorithms that have some degree of adaptability. The down-side, of course, is that these multi-scale refinements can be computationally very expensive, especially when the size of the window is large. The framework of wavelets [53] is a computational efficient approach to multiresolution and has proven to be successful in many applications such as image denoising [109], [110], feature enhancement [111] and shape analysis [112]. In this paper, we use wavelet-related concepts and propose a fast multiresolution wavelet-like algorithm to compute multi-scale local geometric moments of different orders with a dyadic scale progression. In particular, we will consider B-spline window functions, which become wider and more and more Gaussian-like—also meaning isotropic—as the degree of the spline increases.

We believe that these multi-scale local geometric moments could be useful tools for devising new algorithms based on what we call a “sliding window” formulation of a problem. The basic assumption for such an approach is that the spatially-varying feature (or parameter) that one is estimating is approximately constant within the window. The unknown parameter is then estimated from the available information in the window (which often requires the evaluation of moments). Finally, the output value is attributed to the spatial location corresponding to the center of the window. This is a simple, yet powerful paradigm that can be made most effective by working at the appropriate scale (multiresolution strategy). We will illustrate these ideas in Section A.3 by presenting three such local-moment-based algorithms:

- a new method for local shape analysis and feature extraction,
- a multi-scale noise reduction method based on Savitzky-Golay filters [113],
- a multiresolution extension of the Lucas-Kanade optical-flow algorithm [19], which uses a more refined local-affine model for the motion.

These methods are fast thanks to the wavelet-like implementation. The experimental results obtained in all cases are encouraging.

A.2 Theory

In this section, we will define weighted local geometric moments and their associated multiresolution moment filters. We also show how these moments can be computed efficiently in a multiresolution framework.

A.2.1 Weighted Local Geometric Moments

Global geometric moments of order $p \in \mathbb{N}_0$ and location $x_0 \in \mathbb{R}$ of a continuously-defined function f are defined as [96]

$$M_p(x_0) = \int_{\mathbb{R}} (x - x_0)^p f(x) dx. \quad (\text{A.1})$$

For localization, we introduce a positive and symmetric window function w with compact support Ω . We then define weighted local geometric moments of order

p , scale $j \in \mathbb{Z}$ and location x_0 as

$$m_p^{(j)}(x_0) = \int_{\mathbb{R}} (x - x_0)^p w \left(\frac{x - x_0}{2^j} \right) f(x) dx. \quad (\text{A.2})$$

Note that the window function is dilated by a factor 2^j and is centered at x_0 . For a given window function w , we call

$$w_p(x) = x^p w(x) \quad (\text{A.3})$$

the moment filter mask of order p . Then, the local weighted geometric moments can be rewritten in the form of a convolution as

$$m_p^{(j)}(x_0) = 2^{jp} \int_{\mathbb{R}} w_p \left(\frac{x - x_0}{2^j} \right) f(x) dx \quad (\text{A.4})$$

$$= 2^{jp} \left(w_p^{(j)T} * f \right) (x_0), \quad (\text{A.5})$$

where the multiresolution moment filters $w_p^{(j)T}(x) = w_p(-x/2^j)$ are time reversed and dilated versions of the basic moment filter mask (A.3). The normalization factor 2^{jp} in (A.5) is included to simplify the formulation of the multiresolution algorithm presented next.

A.2.2 Two-Scale Equation

Computing local moments at coarser scales becomes more and more time consuming due to the increasing size of the window function. However, multiresolution pyramids of local moments can be computed efficiently, provided that the window function satisfies a two-scale equation, a concept that is closely related to the framework of wavelets [53].

Theorem A.1 (Two-Scale Equation) *Let w be a function that satisfies the two-scale equation*

$$w \left(\frac{x}{2} \right) = \sum_l h(l) w(x - l), \quad (\text{A.6})$$

for some given filter h . Then, w_p satisfies the multi-channel two-scale equation

$$\begin{aligned} 2^p w_p \left(\frac{x}{2} \right) &= \sum_{k=0}^p (h_{p,k} * w_k)(x) \\ &= \sum_{k=0}^p \sum_l h_{p,k}(l) w_k(x - l), \end{aligned} \quad (\text{A.7})$$

with filters $h_{p,k}$, $k = 0, \dots, p$, given by

$$h_{p,k}(l) = \binom{p}{k} l^{p-k} h(l). \tag{A.8}$$

Proof We deduce from (A.3) and (A.6) that

$$w_p\left(\frac{x}{2}\right) = \frac{1}{2^p} \sum_l h(l) x^p w(x-l).$$

Using the fact that

$$x^p = ((x-l) + l)^p = \sum_{k=0}^p \binom{p}{k} l^{p-k} (x-l)^k$$

and applying the definition

$$(x-l)^k w(x-l) = w_k(x-l),$$

we directly obtain (A.7). □

A.2.3 Efficient Multi-Scale Implementation

Theorem A.1 can be used to derive fast algorithms for computing local moments $m_p^{(j)}$ for scales $j = j_0, \dots, j_1$ and orders $p = 0, \dots, P$. To initialize the procedure, the inner products on the finest scale j_0 are computed by using (A.5). Due to Theorem A.1, the coefficients on the subsequent coarser scales can be determined recursively.

Corollary A.2 Let $m_p^{(j)}(n)$, $0 \leq p \leq P$, be local moments at scale j and positions $n \in \mathbb{Z}$. Then, the moments at the next coarser scale $(j+1)$ can be computed as

$$m_p^{(j+1)}(n) = \sum_{k=0}^p \sum_l h_{p,k}^{(j)}(l) m_k^{(j)}(n+2^j l) \tag{A.9}$$

with filter masks $h_{p,k}^{(j)}$ given by

$$h_{p,k}^{(j)} = 2^{j(p-k)} h_{p,k}. \tag{A.10}$$

Proof By definition (A.5) we have that

$$\begin{aligned} m_p^{(j+1)}(n) &= 2^{(j+1)p} \int_{\mathbb{R}} w_p \left(\frac{x-n}{2^{j+1}} \right) f(x) dx \\ &= 2^{(j+1)p} \int_{\mathbb{R}} w_p \left(\frac{1}{2} \frac{x-n}{2^j} \right) f(x) dx. \end{aligned}$$

Using the two-scale equation (A.7) it follows that

$$\begin{aligned} m_p^{(j+1)}(n) &= 2^{(j+1)p} \int_{\mathbb{R}} \frac{1}{2^p} \sum_{k=0}^p \sum_l h_{p,k}(l) w_k \left(\frac{x-n}{2^j} - l \right) f(x) dx \\ &= 2^{jp} \sum_{k=0}^p \sum_l h_{p,k}(l) \int_{\mathbb{R}} w_k \left(\frac{x-(n+2^j l)}{2^j} \right) f(x) dx. \end{aligned}$$

Applying definition (A.5) yields

$$\begin{aligned} m_p^{(j+1)}(n) &= 2^{jp} \sum_{k=0}^p \sum_l h_{p,k}(l) \frac{1}{2^{jk}} m_k^{(j)}(n+2^j l) \\ &= \sum_{k=0}^p \sum_l h_{p,k}(l) 2^{j(p-k)} m_k^{(j)}(n+2^j l). \end{aligned}$$

By defining $h_{p,k}^{(j)}(l) = 2^{j(p-k)} h_{p,k}(l)$, we obtain (A.9). \square

Equation (A.10) means that the two-scale filters $h_{p,k}$ have to be multiplied by $2^{(p-k)}$ at each scale j prior to convolution. The filters $h_{p,k}^{(j)}$ need not to be stored separately since they are obtained by simply updating the basic filters $h_{p,k}$ at each scale. Equation (A.9) is a multi-channel extension of the “à trous” algorithm, which is frequently used for computing overcomplete wavelet transforms [114].

The method is easily modified for computing local moments in a sub-sampled, wavelet-like pyramid. The recursion equation (A.9) then simplifies to a Mallat-like algorithm (cf. [114]):

$$m_p^{(j+1)}(n) = \sum_{k=0}^p \left(h_{p,k}^{(j)T} * m_k^{(j)} \right) (2n), \quad (\text{A.11})$$




Figure A.1: Recursive computation of moments of order 0 to 2 using multi-channel two-scale filters.

where $h_{p,k}^{(j)T}(l) = h_{p,k}^{(j)}(-l)$ denotes the time reversed filter mask $h_{p,k}^{(j)}$. The corresponding block diagram for computing moments of order 0 to 2 in a sub-sampled fashion is shown in Fig. A.1.

To avoid boundary artifacts, the signals that are considered by the algorithm need to be extended properly at the boundary. We assume that the input signal f is extended by using a mirror boundary convention. If this signal is filtered with a symmetric filter (e.g., even moments), the output will exhibit the same symmetry. Conversely, if the signal is filtered with an anti-symmetric filter (e.g., odd moments), the output will be anti-symmetric at the boundary. Therefore, in order to implement the recursive two-scale algorithms (A.9) and (A.11), one has to alternate between the right type of boundary extension of the moments to produce an output that is consistent with the input assumptions. This is ensured by extending even and odd order moments by mirror and anti-mirror boundary conditions, respectively. From (A.8) it can be seen that the two-scale filters $h_{p,k}$ are symmetric or anti-symmetric, if $p - k$ is even or odd, respectively. Thus, the convolution with the properly extended moments $m_k^{(j)}$ will result in the correct boundary extension of the moments $m_p^{(j+1)}$. A summary of all possible cases is given in Table A.1.

The usage of the two-scale algorithm clearly pays off when computing lower order moments at coarser scales. The direct computation of the moments by

Table A.1: Boundary extension of moments.

Order p of $m_p^{(j+1)}$	Bound. extension of $m_p^{(j+1)}$	Order k of $m_k^{(j)}$	Bound. extension of $m_k^{(j)}$	$p - k$	Filter symmetry of $h_{p,k}$
even	mirror	even	mirror	even	symmetric
even	mirror	odd	anti-mirror	odd	anti-symmetric
odd	anti-mirror	even	mirror	odd	anti-symmetric
odd	anti-mirror	odd	anti-mirror	even	symmetric

(A.5) requires $O(2^j)$ multiplications and $O(2^{j+1})$ additions per output point at scale j . On the other hand, the computational complexity of the recursive two-scale algorithm is independent of the scale j and behaves like $O(1)$. A detailed analysis of the computational cost is given in the Appendix A.5.1.

A.2.4 Multiple Dimensions

The notion of multi-scale weighted moments can be extended to multiple dimensions in a straightforward way by using tensor products. In the two-dimensional (2D) case, we define moment filter masks of order $(p + q)$ as

$$w_{p,q}(x, y) = w_p(x)w_q(y). \quad (\text{A.12})$$

The moments at scale j are then given by the separable convolution

$$m_{p,q}^{(j)}(x_0, y_0) = 2^{j(p+q)} \left(w_{p,q}^{(j)T} * f \right) (x_0, y_0), \quad (\text{A.13})$$

where $w_{p,q}^{(j)T}(x, y) = w_p(-x/2^j)w_q(-y/2^j)$ are the associated 2D multiresolution moment filters. For an efficient computation of $m_{p,q}^{(j)}$, equations (A.9) and (A.11) are applied successively in each dimension. In the sub-sampled discrete case, this reads

$$m_{p,q}^{(j+1)}(n, m) = \sum_{k=0}^p \sum_{l=0}^q \left(h_{p,k}^{(j)T} h_{q,l}^{(j)T} * m_{k,l}^{(j)} \right) (2n, 2m), \quad (\text{A.14})$$

where the two-scale filters $h_{p,k}^{(j)T}$ and $h_{q,l}^{(j)T}$ are applied separately in x - and y -directions, respectively. For instance, the block diagram for the second order moment $m_{1,1}^{(j+1)}$ is illustrated in Fig. A.2.




Figure A.2: Recursive computation of $m_{1,1}^{(j+1)}$ using multi-channel two-scale filters.

In the two-dimensional case, the direct moment computation (A.13) requires $O(2^{j+1})$ multiplications and $O(2^{j+2})$ additions per output point at scale j , whereas the cost of the recursive two-scale algorithm behaves as $O(1)$. For a detailed analysis of the computational complexity in two dimensions, we refer to Appendix A.5.1.

A.2.5 B-spline Window Function

The ideal window function should be positive, with weights decreasing away from the center, refinable, separable, and isotropic in multiple dimensions. The only choice would be a Gaussian, but it does not satisfy a two-scale equation. However, B-splines β^n satisfy a two-scale equation and rapidly converge to Gaussians when their degree $n \in \mathbb{N}$ increases [51]. In fact, for a given number of filter tabs, B-splines are the smoothest scaling functions in the Sobolev sense [115]; this guaranties that they converge fastest to Gaussians in the Sobolev norm. This ensures nearly isotropy of the window in multiple dimensions. The cubic B-spline ($n = 3$), β^3 , and its two first moment filters β_1^3 and β_2^3 are plotted in Fig. A.3. The corresponding two-scale filters $h_{p,k}$ up to order $p = 2$ are given in Table A.2.

The Fourier transform of a B-spline β^n , which is the $(n + 1)$ -fold convolution of a rectangular pulse, is given by

$$\widehat{\beta}^n(\omega) = \left(\frac{\sin(\omega/2)}{\omega/2} \right)^{n+1}. \tag{A.15}$$




Figure A.3: Cubic B-spline β^3 and its first two moment filters β_1^3 and β_2^3 .

Table A.2: Two-scale filters $h_{p,k}$ up to order $p = 2$ for β^3 .

l	-2	-1	0	1	2
$h_{0,0}(l)$	$\frac{1}{8}$	$\frac{4}{8}$	$\frac{6}{8}$	$\frac{4}{8}$	$\frac{1}{8}$
$h_{1,0}(l)$	$-\frac{1}{4}$	$-\frac{2}{4}$	0	$\frac{2}{4}$	$\frac{1}{4}$
$h_{1,1}(l)$	$\frac{1}{8}$	$\frac{4}{8}$	$\frac{6}{8}$	$\frac{4}{8}$	$\frac{1}{8}$
$h_{2,0}(l)$	$\frac{1}{2}$	$\frac{1}{2}$	0	$\frac{1}{2}$	$\frac{1}{2}$
$h_{2,1}(l)$	$-\frac{1}{2}$	$-\frac{2}{2}$	0	$\frac{2}{2}$	$\frac{1}{2}$
$h_{2,2}(l)$	$\frac{1}{8}$	$\frac{4}{8}$	$\frac{6}{8}$	$\frac{4}{8}$	$\frac{1}{8}$

By definition, the Fourier transforms of the corresponding moment filters are given by

$$\widehat{\beta}_p^n(\omega) = i^p \frac{d^p}{d\omega^p} \widehat{\beta}^n(\omega). \quad (\text{A.16})$$

B-splines of degree n are by construction in C^{n-1} , i.e., they are $(n-1)$ times continuously differentiable; the same also holds true for the moment filters. This implies that their Fourier transforms decay at least like $O(1/|\omega|^n)$ for large ω . Consequently, the Fourier transforms of the moment filters decay faster when the spline degree increases. Fig. A.4 shows the normalized spectra of the B-spline $\beta^5(x/2)$ and its moment filters $\beta_p^5(x/2)$ for degree $n = 5$ and moment orders $p = 0, 1, 2$ at scale $j = 1$. It is clear from this graph that the filters are essentially bandpass, which can be used as a justification for the downsampling




Figure A.4: Normalized spectra of moment filters $\beta^5(x/2)$, $\beta_1^5(x/2)$ and $\beta_2^5(x/2)$.

of moments at coarser scales.

A.3 Applications

The fast algorithm presented above is applicable to a variety of image analysis problems, such as image segmentation, pattern detection, and optical-flow estimation, for which local solutions over sliding windows have been proposed. These approaches can be extended by applying a multiresolution strategy which provides adaptability while also reducing computational cost. Here, we will illustrate the concept by presenting new local-moment-based algorithms for three specific tasks: (1) local shape analysis and feature extraction, (2) filtering for noise reduction, and (3) the estimation of motion fields using a local affine model.

A.3.1 Local Shape Analysis and Feature Extraction

Effective analysis of shapes is required by many computer vision applications; in particular, in biomedical image analysis. One of the major issues is to determine location, orientation and size features of filamentous or spherical bright structures in an image. Examples are segmentation and characterization of biological cell images, the analysis of vessel distributions in medical images and the detection of DNA filaments in electron micrograph images. The evaluation of low order moments represents a systematic and efficient method of shape analysis. Since moments are integral-based features, they are robust against noise.

Furthermore, low order moments have a direct geometrical interpretation.

Geometric Interpretation of Moments

The moments $m_{p,q}$ have well-defined geometric interpretations. The coordinates of the local centroid are given by

$$\bar{x} = m_{1,0}/m_{0,0} \quad \text{and} \quad \bar{y} = m_{0,1}/m_{0,0}. \quad (\text{A.17})$$

The distance between the window center and the local centroid allows to detect whether the sliding window is located on the center of a bright structure or not. The so-called central moments [96] can be expressed in terms of ordinary moments $m_{p,q}$ and the coordinates of the centroid. For the second order we have

$$\mu_{2,0} = m_{2,0} - m_{0,0}\bar{x}^2, \quad \mu_{0,2} = m_{0,2} - m_{0,0}\bar{y}^2, \quad (\text{A.18})$$

$$\mu_{1,1} = m_{1,1} - m_{0,0}\bar{x}\bar{y}. \quad (\text{A.19})$$

These three central moments of second order are the components of the inertia matrix

$$\mathbf{J} = \begin{pmatrix} \mu_{2,0} & \mu_{1,1} \\ \mu_{1,1} & \mu_{0,2} \end{pmatrix}. \quad (\text{A.20})$$

The local orientation of the analyzed object is given by the eigenvector corresponding to the minimal eigenvalue of \mathbf{J} . In fact, the local object is mapped onto an ellipsoid centered at (\bar{x}, \bar{y}) . The ellipsoid axes are directed along the eigenvectors of \mathbf{J} and the corresponding axes semi-lengths are the magnitudes of the respective eigenvalues λ_1 and λ_2 . The orientation angle with respect to the x -axis is given by

$$\phi = \frac{1}{2} \arctan \left(\frac{2\mu_{1,1}}{\mu_{2,0} - \mu_{0,2}} \right). \quad (\text{A.21})$$

A measure for the eccentricity of the local ellipsoid is given by

$$\varepsilon = \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \right)^2 = \frac{(\mu_{2,0} - \mu_{0,2})^2 + 4\mu_{1,1}^2}{(\mu_{2,0} + \mu_{0,2})^2} \quad (\text{A.22})$$

and takes values between 0 and 1. It indicates whether the local object is elongated or not. The eccentricity measure is independent from the local image region and is therefore well-suited for inter-scale comparisons.

Multi-Scale Detection Strategy

Brighter elongated structures or filaments can be extracted by evaluating the various moment features and putting thresholds on eigenvalues or eccentricity measures. Since the elongated structures of interest can have different sizes, we propose to detect them at multiple scales $j_0 \leq j \leq j_1$, where j_0 and j_1 are the finest and coarsest scale at which relevant structures are expected. A simple strategy, which was applied in our experiments, is described in the following: At each image pixel (n_x, n_y) we compute the local moments $m_{p,q}^{(j)}$ for $(p+q) \leq 2$. From these we derive the local orientations and eccentricities $\varepsilon^{(j)}$. To decide whether or not a local object is part of a filamentous structure, we compute the figure of merit

$$\gamma^{(j)} = \varepsilon^{(j)} e^{-(\bar{x}^2 + \bar{y}^2)/(2^{(2j+1)}\sigma^2)}. \quad (\text{A.23})$$

The second factor in (A.23) assigns more weight to cases where the local centroid (\bar{x}, \bar{y}) is close to the center of the local window. The parameter σ controls the range of the centroid around the window origin to be accepted. The multi-scale approach also helps us to detect cases where the local structure is located symmetrically at the periphery of the window function. To avoid these cases, the figure of merit $\gamma^{(j)}$ is set to zero, if $m_{0,0}^{(j-1)} < m_{0,0}^{(j)}$. This means that the local mean of the gray values at the next finer scale has to be greater than the local mean at the current scale.

The figure of merit (A.23) will be maximal at a scale that approximately matches the size of the elongated shape to detect. Therefore, we integrate the figures of merit at different scales to obtain a final estimate for the goodness of local fit by

$$\gamma = \max_{j_0 \leq j \leq j_1} \gamma^{(j)}. \quad (\text{A.24})$$

Application: Detection of DNA Filaments

The structure of DNA molecules can be visualized by cryo-electron-microscopy (CEM) [116]. Because of the physical process involved, the resulting images have very low contrast to avoid destruction of the specimen (cf. Fig. A.6). Biologists are highly interested in an automatic detection of the thin strands of DNA, but the task is challenging because of the poor signal-to-noise ratio (near 0 dB).

The proposed moment-based algorithm was tested on synthetic and real images. Fig. A.5a and A.5b show a synthetic circular DNA strand with two

different levels of additive Gaussian noise, respectively. In this experiment, we used a B-spline window of degree 3 at scales $j = 2, 3$.


The algorithm is compared with the so-called structure tensor method [67], a standard method to estimate local orientations of image patterns. Instead of the inertia tensor (A.20), this method uses the structure tensor

$$\mathbf{S} = \begin{pmatrix} \langle w, I_x^2 \rangle & \langle w, I_x I_y \rangle \\ \langle w, I_x I_y \rangle & \langle w, I_y^2 \rangle \end{pmatrix}, \quad (\text{A.25})$$


where $w(x, y)$ denotes a window function and I_x, I_y denote the partial derivatives of the image intensity $I(x, y)$. The computation of the local orientation and eccentricity measure is analog to (A.21) and (A.22), respectively. As in [67], we interpret the estimated eccentricity as a figure of merit and use a Gaussian window function. For the standard deviation of the Gaussian window, we used $\sigma = 1.7$ which corresponds to the effective width of the B-spline window at the finest scale $j = 2$ of the moment-based algorithm.

The estimated eccentricities of the structure tensor approach are shown in Fig. A.5c and A.5d for the two different noise levels, respectively. Fig. A.5e and A.5f show the corresponding figures of merit of the proposed moment-based algorithm. For the lower noise level both methods detect the circular structure well. However, the figure of merit of the moment-based algorithm is much thinner around the true structure since the eccentricity measure is weighted by the distance of the window center to the centroid of the local image content as described in (A.23). This feature is not available in the structure tensor approach. In the case of the higher noise level, the moment based algorithm still detects the elongated object fairly well (Fig. A.5f). In contrast, the structure tensor approach degrades significantly (Fig. A.5d). This is probably due to the fact that this method uses derivatives which are sensitive to noise, whereas the proposed approach is integral-based.


The moment-based detection algorithm was also applied to real images as shown in Fig. A.6a. Since the intensity in CEM-images may vary globally, the original images were first normalized in a pre-processing step. We used moments of order zero (local average) at scale $j = 2$ for local background subtraction. Then we computed for each pixel the figure of merit γ as described above. In particular, we used a B-spline window of degree 3 at scales $j = 2, 3$. The figures of merit were then thresholded to suppress values that correspond to non-significant structures. The final figures of merit are visualized in Fig. A.6b in form of a needle diagram. The length of the needles is proportional to the size of the figure of merit at each pixel. The direction of the needles corresponds to




(a) Synthetic elongated structure with additive Gaussian noise (SNR = 28.14 dB)



(b) Synthetic elongated structure with additive Gaussian noise (SNR = 8.15 dB)




(c) Figure of merit of structure tensor-based algorithm applied to (a)




(d) Figure of merit of structure tensor-based algorithm applied to (b)

Figure A.5: Comparison of moment-based and inertia tensor-based detection of elongated structures for different noise levels.




(e) Figure of merit of moment-based algorithm applied to (a)




(f) Figure of merit of moment-based algorithm applied to (b)

Figure A.5: Comparison of moment-based and inertia tensor-based detection of elongated structures for different noise levels.



(a) Original CEM-image



(b) Estimated local orientation

Figure A.6: A CEM-image and detected DNA strands.

the local orientation of the object. We see that the two DNA strands contained in the image together with their local orientation were clearly detected. Failures due to the high noise content in the image are very sparse.

A.3.2 Multi-Scale, Weighted Savitzky-Golay Smoothing Filters

Savitzky-Golay filtering [113] can be thought of as a generalized moving average filter. The idea of Savitzky-Golay filtering is to find filter coefficients that preserve higher order polynomials. These filter coefficients are derived by a least-squares fitting of a polynomial of given degree within a sliding window. The smoothed points are computed by replacing each data point with the value of the fitted polynomial at the window center. For this reason, a Savitzky-Golay filter is also called a digital smoothing polynomial filter or a least-squares smoothing filter. A crucial point is the choice of the size of the window function. A small window preserves narrow features of the underlying signal, but filters less; larger windows smooth more, but lead to blurring of image details.

Originally, this approach was proposed for one-dimensional signals and used a box-shaped window function of fixed length [113]. Here, we propose a multi-dimensional extension based on a weighted least squares criterion. We also propose a new multi-scale filtering strategy whereby the final smoothed image is obtained by combining results from different scales using a hypothesis test.

Weighted Savitzky-Golay Filtering

Let us consider a two-dimensional polynomial of degree d

$$P_d(x, y) = \sum_{0 \leq p+q \leq d} a_{p,q} x^p y^q, \quad (\text{A.26})$$

which is specified by the $N_a = (d+1)(d+2)/2$ polynomial coefficients $a_{p,q}$. Let w denote a window function with discrete support Ω of cardinality N_Ω which is located at (m_0, n_0) . To fit the polynomial locally to an image $I(m, n)$, we minimize the weighted least-squares functional

$$r^2 = \sum_{(m,n) \in \Omega} w(m, n) \left(P_d(m, n) - I(m + m_0, n + n_0) \right)^2. \quad (\text{A.27})$$

By differentiating (A.27) with respect to each of the unknown polynomial coefficients $a_{p,q}$, we obtain the corresponding normal equations $\mathbf{A}^T \mathbf{W} \mathbf{A} \mathbf{a} = \mathbf{A}^T \mathbf{W} \mathbf{b}$,

where

$$\{\mathbf{A}^T \mathbf{W} \mathbf{A}\}_{p_1, q_1; p_2, q_2} = \sum_{(m, n) \in \Omega} m^{p_1 + p_2} n^{q_1 + q_2} w(m, n), \quad (\text{A.28})$$

$$\{\mathbf{A}^T \mathbf{W} \mathbf{b}\}_{p_1, q_1} = \sum_{(m, n) \in \Omega} m^{p_1} n^{q_1} w(m, n). \quad (\text{A.29})$$

$$I(m + m_0, n + n_0),$$

and

$$\{\mathbf{a}\}_{p_2, q_2} = a_{p_2, q_2}, \quad (\text{A.30})$$

with $0 \leq p_1 + q_1 \leq d$ and $0 \leq p_2 + q_2 \leq d$. The index-tuples (p_1, q_1) and (p_2, q_2) denote the row and column indices of the matrix and vectors, respectively. The $N_\Omega \times N_\Omega$ diagonal matrix \mathbf{W} is composed by the weights $w(m, n)$. Since $\mathbf{A}^T \mathbf{W} \mathbf{A}$ does not depend on the image data, the matrix and its inverse can be computed once and forever in advance. The right hand side vector $\mathbf{A}^T \mathbf{W} \mathbf{b}$ is nothing but a discrete version of the local moments (A.13) of order zero to d . The smoothed image point at the window center is equal to the polynomial coefficient $a_{0,0}$, which is given by the inner product of the corresponding row $\{(\mathbf{A}^T \mathbf{W} \mathbf{A})^{-1}\}_{0,0}$ of the matrix inverse and the right hand side $\mathbf{A}^T \mathbf{W} \mathbf{b}$.

Multi-Scale Strategy

In order to find a trade-off between the conflicting requirements of noise reduction and conservation of image details, we propose a multi-scale framework of the introduced weighted Savitzky-Golay filtering. We assume that the image is locally given by the model: polynomial signal + noise, i.e.,

$$I = P_d + \varepsilon, \quad (\text{A.31})$$

where $\varepsilon \sim N(0, \sigma^2)$ corresponds to Gaussian white noise of zero mean and common variance σ^2 . Thus, the residual (A.27) gives the expected squared deviation of the image data from the given polynomial P_d due to noise. As shown in the appendix, the normalized residual corresponds to a linear combination of $N_\Omega - N_a$ independent χ_1^2 -distributed random variables, i.e.,

$$\frac{r^2}{\sigma^2} \sim \sum_{n=1}^{N_\Omega - N_a} \lambda(n) \chi_1^2, \quad (\text{A.32})$$

where the coefficients $\lambda(n)$ are given by the $N_\Omega - N_a$ non-zero eigenvalues of the matrix

$$\mathbf{W} - \mathbf{W}\mathbf{A}(\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^T\mathbf{W}. \quad (\text{A.33})$$

For a proof and the computation of the probability density function of (A.32) we refer to Appendix A.5.2. Note that, for uniform weights, the resulting distribution (A.32) corresponds to a $\chi^2_{N_\Omega - N_a}$ distribution with $N_\Omega - N_a$ degrees of freedom. Asymptotically, r^2/σ^2 is normally distributed with mean

$$E\left(\frac{r^2}{\sigma^2}\right) = \sum_{n=1}^{N_\Omega - N_a} \lambda(n) \quad (\text{A.34})$$

and variance

$$\text{var}\left(\frac{r^2}{\sigma^2}\right) = 2 \sum_{n=1}^{N_\Omega - N_a} \lambda(n)^2. \quad (\text{A.35})$$

When working on real images, (A.32) enables us to detect image regions for which the chosen polynomial degree or window size are not adequate. More specifically, we apply a two-sided hypothesis test on r^2/σ^2 with a given significance level α . In order to avoid cases where the degree of the polynomial is too high for the given image structure and tends to fit the noise, we reject results for which the residual is below the confidence interval. This usually happens when using small windows in flat image regions. On the other hand, we also reject results for which the residual is above the confidence interval. In this case, image details like edges cannot be fitted closely by the polynomial. The aim is to use locally a window as large as possible to achieve maximum noise reduction. Consequently, we compute smoothed image versions using windows at scales $j = j_0, \dots, j_1$. Recall that the images of moments (A.29) can be computed efficiently for different scales by using (A.9). The final smoothed image is obtained by choosing, for each pixel, the output value from the coarsest scale for which the normalized residual remains inside the confidence interval.

Numerical Results

In order to demonstrate the performance of weighting and multi-scale filtering, we have applied the algorithm to an image containing additive Gaussian white noise. Fig. A.7a shows the original image and Fig. A.7b shows the image after adding Gaussian white noise of standard deviation $\sigma = 20.0$, resulting in a




Figure A.7: Original and noisy image.

signal-to-noise ratio of 20.40 dB. Results were computed for a B-spline window of degree 3 and a fitting polynomial of degree 2.


Fig. A.8 illustrates the effect of using B-spline weighting. The left column displays the filtered outputs of B-spline-weighted Savitzky-Golay filtering at scales $j = 1, \dots, 3$. Window sizes were 7, 15 and 31 pixels at each scale, respectively. The right column corresponds to the case of using a squared window with constant weights one. Here, the support was chosen to be the effective duration of the B-spline windows, resulting in window sizes of 3, 5 and 9 pixels, respectively. Fig. A.9a shows the final output image of the multi-scale, B-spline-weighted method. Smoothed image versions at scales $j = 1, \dots, 3$ were combined to the final output image using a double-sided hypothesis test on the normalized residuals (A.32) with a significance level $\alpha = 0.01$. The signal-to-noise ratio of the final image is 27.30 dB, which is significantly larger than the signal-to-noise ratios at the single scales (26.25 dB, 24.55 dB, 21.63 dB for scales $j = 1, \dots, 3$). Also visually, the final output image seems to be superior to the single scale outputs. Image details like edges are well preserved, whereas flat image regions are fairly smoothed. The result is compared with two standard denoising algorithms. The first is a wavelet soft-thresholding method. The noisy image was decomposed in a 3-level wavelet transform pyramid using orthogonal Battle-Lemarié wavelets [117]. We used the same order of spline ($n = 3$) for the methods to be comparable. We also optimized the method by selecting the threshold $T = 26$, yielding the maximum signal-to-noise ratio (SNR = 24.89 dB). From Fig. A.9c it can be seen that the wavelet-based smoothed image is clearly more blurred and suffers from typical ringing artifacts. The second com-




(a) B-spline-weighted at scale $j = 1$ (SNR = 26.25 dB)



(b) Non-weighted at scale $j = 1$ (SNR = 22.67 dB)



(c) B-spline-weighted at scale $j = 2$ (SNR = 24.55 dB)



(d) Non-weighted at scale $j = 2$ (SNR = 25.49 dB)

Figure A.8: Images (a) to (d) demonstrate the effect of B-spline weighting at scales $j = 1, 2$.




Figure A.8: Images (e) and (f) demonstrate the effect of B-spline weighting at scale $j = 3$.

parison method is the adaptive Wiener filter [118]. This filter corresponds to a pixel-wise adaptive Wiener method based on statistics derived from a local neighborhood of each pixel. The maximum signal-to-noise ratio of 25.96 dB was obtained for a filter size of (5×5) pixels. As can be seen from Fig. A.9d, the Wiener filter preserves image details well, but smooths less in flat image regions.


In the present approach, different scales are combined in an exclusive fashion which leads to some artifacts near edge regions. Although the proposed multi-scale denoising algorithm performs best in terms of signal-to-noise-ratio, it may be possible to improve the visual perception of the output further by using a more progressive weighted combination.

A.3.3 Optical-Flow Estimation

The estimation of motion from an image sequence is a classical problem in computer vision. Among others, the optical-flow technique has been proven to be a successful approach to this problem [61].

Let $I(x, y, t)$ denote the intensity of pixels at location $\mathbf{r} = (x, y)$ and time t in an image sequence. Gradient-based optical flow estimation relies on the assumption that the intensity of a particular point in a moving pattern does not change with time. The constant intensity assumption can be expressed as [16]


$$I_x(\mathbf{r}, t) u(\mathbf{r}, t) + I_y(\mathbf{r}, t) v(\mathbf{r}, t) + I_t(\mathbf{r}, t) = 0.$$




(a) Final multi-scale, B-spline-weighted Savitzky-Golay output (SNR = 27.30 dB)



(b) Final multi-scale, non-weighted Savitzky-Golay output (SNR = 26.42 dB)



(c) Wavelet thresholded image (SNR = 24.89 dB)



(d) Wiener filtered image (SNR = 25.96 dB)

Figure A.9: Comparison of different smoothing methods.

I_x , I_y and I_t denote the spatial and temporal derivatives of the image intensity. The velocities u and v are, respectively, the x - and y -components of the optical-flow we wish to estimate.

Local Affine Motion

A very popular optical-flow algorithm is the Lucas-Kanade method [19], which estimates the motion locally, assuming the motion to be constant within a window of support Ω . In order to account for more complex motions, such as rotation, divergence, and shear, we extend this approach to a local affine model for the motion. If (x_0, y_0) denotes the center of the local window, this model is defined as

$$\begin{aligned} u(x, y) &= u_0 + u_x(x - x_0) + u_y(y - y_0), \\ v(x, y) &= v_0 + v_x(x - x_0) + v_y(y - y_0). \end{aligned} \tag{A.36}$$

The parameters u_0 and v_0 correspond to the motion at the window center and u_x , u_y , v_x , and v_y are the first order spatial derivatives of u and v , respectively. The local motion components can be estimated by minimizing the weighted least-squares criterion

$$\int_{\mathbb{R}^2} w(x - x_0, y - y_0) \left(I_x u + I_y v + I_t \right)^2 dx dy. \tag{A.37}$$

The symmetric window function w gives more weight to constraints at the center of the local region than to those at the periphery. By differentiating (A.37) with respect to each of the six unknown parameters, we obtain the so-called normal equations $\mathbf{A}^T \mathbf{W} \mathbf{A} \mathbf{v} = \mathbf{A}^T \mathbf{W} \mathbf{b}$ in terms of local moments of orders zero to two of the spatial and temporal derivatives of I as defined in (A.38) on page 143.

Coarse-To-Fine Multi-Scale Strategy

It is obviously difficult to estimate large motions at fine scales. A way around this problem is to apply a coarse-to-fine strategy. At each scale $j_0 \leq j \leq j_1$, we compute the local moments on a grid which is sub-sampled by 2^j in each dimension. These sub-sampled, multi-scale local moments can be computed efficiently by using (A.11).

The motion vectors are cascaded through each resolution level as initial estimates and are then replaced if they do not already exceed a scale-dependent

$$\begin{aligned}
& \mathbf{A}^T \mathbf{W} \mathbf{A} = \\
& \left(\begin{array}{cccccc}
\langle w, I_x^2 \rangle & \langle w, I_x I_y \rangle & \langle x w, I_x^2 \rangle & \langle y w, I_x^2 \rangle & \langle x w, I_x I_y \rangle & \langle y w, I_x I_y \rangle \\
\langle w, I_x I_y \rangle & \langle w, I_y^2 \rangle & \langle x w, I_x I_y \rangle & \langle y w, I_x I_y \rangle & \langle x w, I_y^2 \rangle & \langle y w, I_y^2 \rangle \\
\langle x w, I_x^2 \rangle & \langle x w, I_x I_y \rangle & \langle x^2 w, I_x^2 \rangle & \langle x y w, I_x^2 \rangle & \langle x^2 w, I_x I_y \rangle & \langle x y w, I_x I_y \rangle \\
\langle y w, I_x^2 \rangle & \langle y w, I_x I_y \rangle & \langle x y w, I_x^2 \rangle & \langle y^2 w, I_x^2 \rangle & \langle x y w, I_x I_y \rangle & \langle y^2 w, I_x I_y \rangle \\
\langle x w, I_x I_y \rangle & \langle x w, I_y^2 \rangle & \langle x^2 w, I_x I_y \rangle & \langle x y w, I_x I_y \rangle & \langle x^2 w, I_y^2 \rangle & \langle x y w, I_y^2 \rangle \\
\langle y w, I_x I_y \rangle & \langle y w, I_y^2 \rangle & \langle x y w, I_x I_y \rangle & \langle y^2 w, I_x I_y \rangle & \langle x y w, I_y^2 \rangle & \langle y^2 w, I_y^2 \rangle
\end{array} \right) \\
& \mathbf{v} = \begin{pmatrix} u_0 \\ v_0 \\ u_x \\ u_y \\ v_x \\ v_y \end{pmatrix}, \quad \mathbf{A}^T \mathbf{W} \mathbf{b} = - \begin{pmatrix} \langle w, I_x I_t \rangle \\ \langle w, I_y I_t \rangle \\ \langle x w, I_x I_t \rangle \\ \langle y w, I_x I_t \rangle \\ \langle x w, I_y I_t \rangle \\ \langle y w, I_y I_t \rangle \end{pmatrix} \quad (\text{A.38})
\end{aligned}$$

size. For each local estimate, we compute the confidence measure

$$1 - \sin \theta = 1 - \frac{\|\mathbf{W}^{1/2} (\mathbf{A} \mathbf{v} - \mathbf{b})\|_{l_2}}{\|\mathbf{W}^{1/2} \mathbf{b}\|_{l_2}} \in [0, 1]. \quad (\text{A.39})$$

The argument θ corresponds to the angle between the vectors $\mathbf{W}^{1/2} \mathbf{b}$ and $\mathbf{W}^{1/2} \mathbf{A} \mathbf{v}$ and characterizes how close $\mathbf{W}^{1/2} \mathbf{b}$ is to the image of $\mathbf{W}^{1/2} \mathbf{A}$. A local estimate is replaced only if its confidence measure is larger than the corresponding one at the next coarser scale. Otherwise, the coarser scale estimator is kept. Furthermore, a solution of a local linear system is regarded as not admissible if the linear system is either ill-conditioned or if the length of the estimated central motion vector exceeds some scale-dependent limit. Finally, a motion estimate is set to zero if the local mean of the time derivative at the given location is below a pre-defined noise level.

The final motion estimates at the finest scale j_0 are then interpolated by B-splines to obtain a continuous representation of the motion field.

Numerical Results

The performance of the algorithm was tested on synthetic and real image sequences. In particular, we used the well known synthetic sequence ‘‘Yosemite’’. Since the exact motion field is known, the error of the estimated motion field was computed using the angular error measure as defined in [61]. As real data




Figure A.10: One frame of the Yosemite sequence and its corresponding estimated motion field.

we used the “Rubik Cube” sequence¹. One frame of each sequence and its corresponding estimated motion field are shown in Fig. A.10 and A.11. All sequences were prefiltered with a Binomial filter of variance $\sigma^2 = 1.5$ and a B-spline window of degree 5 at scales $j = 2, \dots, 4$ was used for moment computation.

The angular error of the “Yosemite” sequence is $6.33^\circ \pm 9.98^\circ$ with a flow field density of 100%. The error of the corresponding adaptation of the Lucas-Kanade approach (same window, same multiresolution strategy, locally constant motion model) is $7.43^\circ \pm 12.72^\circ$. Barron & al. [61] report an average angular error of an optimized Horn and Schunk method [16] (spatio-temporal prefiltering, 4-point central differences for differentiation) of $11.26^\circ \pm 16.41^\circ$ with a flow field density of 100%. Their implementation of an improved version of the original Lucas-Kanade method (spatio-temporal prefiltering, rejecting unreliable estimates) only produced a reasonable error for a very sparse velocity field with a density of 35.1%.

The rotational movement in the “Rubik Cube” sequence is also clearly recovered. The obtained results also compare favorably with all other methods evaluated in the survey of Barron & al.

¹All sequences were downloaded from Barron & al.’s FTP site at <ftp://csd.uwo.ca/pub/vision>.

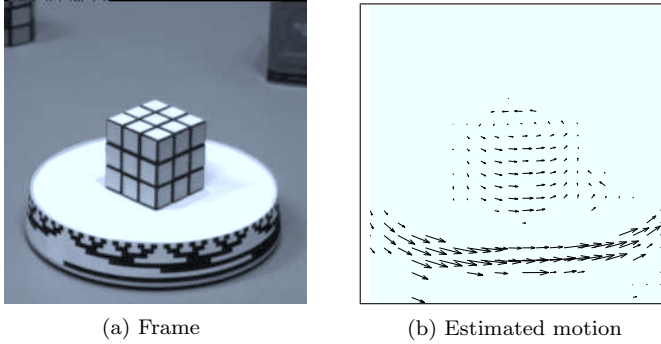


Figure A.11: One frame of the Rubik Cube sequence and its corresponding estimated motion field.

A.4 Conclusions

We have introduced B-spline-weighted, local geometric moments within windows of dyadic sizes. The weighting ensures isotropy in multiple dimensions and the scalability allows adaptability to local image contents. Computational efficiency was achieved by developing a Mallat-like algorithm to compute these moments at multiple scales.

Local moments provide a powerful set of features that can be used in many sliding-window-type algorithms. In particular, we demonstrated their usefulness on three different image analysis problems: feature extraction, noise reduction, and optical-flow estimation. We proposed basic, moment-based algorithms with promising experimental results. Some aspects of these generic algorithms can be further improved by tuning them to special applications. Besides the applications mentioned, these moments could also be useful for applications such as pattern classification and image segmentation.

A.5 Appendix

A.5.1 Computational Complexity

In the following, we analyze the computational complexity of the recursive two-scale algorithm in the one and two-dimensional case.





Figure A.12: 1D computational complexity per output point for direct moment computation and usage of two-scale algorithm.

Computational Complexity in 1D

We assume that the length of the discretized window function w is $2N + 1$ and has a corresponding two-scale filter of length $2N + 3$. Since the window function is symmetric, the direct calculation of (A.5) requires $2^j(N + 1)$ multiplications and $2^{j+1}(N + 1) - 2$ additions per output point at scale j , independently of the order p . On the other hand, since the two-scale filters are either symmetric or anti-symmetric, the two-scale algorithms (A.9) and (A.11) require $(p + 1)(N + 2)$ multiplications and $2(p + 1)(N + 1) + p$ additions for moment order p and are independent of the scale. The computational complexities for scales $j = 1, \dots, 5$, moment orders $p = 0, 1, 2$ and $N = 1$ are plotted in Fig. A.12. Obviously, the use of the two-scale algorithm starts paying off at scales $j = 1, 2, 3$ for moment orders $p = 0, 1, 2$, respectively. Since in practice usually low order moments are used, the proposed computation scheme is much more efficient at coarser scales.

Computational Complexity in 2D

For the two-scale algorithm (A.14), the number of multiplications and additions per output point at scale j are $2(p + 1)(q + 1)(N + 2)$ and $4(p + 1)(q + 1)(N + 1) + (p + 1)(q + 1) - 1$, respectively; in contrast, the direct computation (A.13) requires $2^{j+1}(N + 1)$ multiplications and $2^{j+2}(N + 1) - 4$ additions. The computational complexities for scales $j = 1, \dots, 5$, moment orders $0 \leq p + q \leq 2$ and $N = 1$ are plotted in Fig. A.13. The two-scale algorithm clearly pays off at coarser scales.



(a) Number of multiplications

(b) Number of additions

Figure A.13: 2D computational complexity per output point for direct moment computation and usage of two-scale algorithm.

A.5.2 Computation of the Probability Density Function of a Weighted Least-Squares Residual

Let $\mathbf{Ax} = \widehat{\mathbf{b}}$ be an overdetermined linear system of size $n \times p$, $n > p$, and maximum rank p . We assume that the noisy observation $\widehat{\mathbf{b}}$ is given by $\widehat{\mathbf{b}} = \mathbf{b} + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$ is jointly normally distributed and $\mathbf{b} = \mathbf{Ax}$.

The weighted least-squares estimator \mathbf{x} is obtained by minimizing

$$r^2 = (\mathbf{Ax} - \widehat{\mathbf{b}})^T \mathbf{W} (\mathbf{Ax} - \widehat{\mathbf{b}}),$$

where \mathbf{W} is a diagonal $(n \times n)$ -matrix of weights. Using the fact that $\mathbf{x} = (\mathbf{A}^T \mathbf{WA})^{-1} \mathbf{A}^T \mathbf{W} \widehat{\mathbf{b}}$ and that $\widehat{\mathbf{b}} = \mathbf{Ax} + \boldsymbol{\varepsilon}$, we obtain

$$r^2 = \boldsymbol{\varepsilon}^T \mathbf{C} \boldsymbol{\varepsilon}, \quad (\text{A.40})$$

where $\mathbf{C} = \mathbf{W} - \mathbf{WA}(\mathbf{A}^T \mathbf{WA})^{-1} \mathbf{A}^T \mathbf{W}$. Since \mathbf{C} is symmetric, it can be decomposed as $\mathbf{C} = \mathbf{U}^T \boldsymbol{\Lambda} \mathbf{U}$, where \mathbf{U} is an orthogonal matrix and $\boldsymbol{\Lambda}$ is a real diagonal matrix containing the eigenvalues of \mathbf{C} . Therefore, (A.40) can be expressed as

$$\begin{aligned} r^2 &= (\mathbf{U}\boldsymbol{\varepsilon})^T \boldsymbol{\Lambda} (\mathbf{U}\boldsymbol{\varepsilon}) \\ &= \boldsymbol{\eta}^T \boldsymbol{\Lambda} \boldsymbol{\eta}, \end{aligned}$$

where $\boldsymbol{\eta} = \mathbf{U}\boldsymbol{\varepsilon}$ is also a $N(\mathbf{0}, \sigma^2 \mathbf{I})$ -distributed random variable due to the orthogonality of \mathbf{U} . Since \mathbf{C} is by construction of rank $(n - p)$, we have that

$$r^2 = \sum_{k=1}^{n-p} \lambda(k) \boldsymbol{\eta}^2(k),$$

where $\lambda(k)$ denote the non-zero diagonal elements of $\boldsymbol{\Lambda}$.

Now, the $\boldsymbol{\eta}(k)/\sigma \sim N(0, 1)$ are independently normally distributed so that their squares follow a χ_1^2 -distribution. Consequently, the probability density function (pdf) of r^2/σ^2 is given by the convolution of χ_1^2 -pdf's dilated and scaled by the factors $\lambda(k)$. Since the characteristic function (Fourier transform of the pdf) of a χ_1^2 -distribution is given by

$$f_{\chi_1^2}(\omega) = \frac{1}{(1 - 2i\omega)^{1/2}},$$

the characteristic function of r^2/σ^2 is given by

$$f_{r^2/\sigma^2}(\omega) = \prod_{k=1}^{n-p} \frac{1}{(1 - 2i\lambda(k)\omega)^{1/2}}.$$

For n sufficiently large, the pdf converges to a Gaussian as a consequence of the central limit theorem.

Bibliography

- [1] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: active contour models," *Intl. J. Computer Vision*, vol. 1, pp. 321–332, Jan. 1988.
- [2] V. Chalana, D. Linker, D. Haynor, and Y. Kim, "A multiple active contour model for cardiac boundary detection on echocardiographic sequences," *IEEE Transactions on Medical Imaging*, vol. 15, no. 3, pp. 290–298, June 1996.
- [3] J. McEachen II and J. Duncan, "Shape-based tracking of left ventricular wall motion," *IEEE Transactions on Medical Imaging*, vol. 16, no. 3, pp. 270–283, June 1997.
- [4] T. Hozumi, K. Yoshida, H. Yoshioka, T. Yagi, T. Akasaka, T. Takagi, M. Nishiura, M. Watanabe, and J. Yoshikawa, "Echocardiographic estimation of left ventricular cavity area with a newly developed automated contour tracking method," *Journal of the American Society of Echocardiography*, vol. 10, no. 8, pp. 822–829, Oct. 1997.
- [5] T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models - their training and application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, Jan. 1995.
- [6] A. Parker, A. Hill, C. Taylor, T. Cootes, X. Jin, and D. Gibson, "Application of point distribution models to the automated analysis of echocardiograms," in *Proceedings of Computers in Cardiology*. Bethesda, MD, USA: IEEE Computer Society Press, September, 1994, pp. 25–28.
- [7] S. Malassiotis and M. Strintzis, "Tracking the left ventricle in echocardiographic images by learning heart dynamics," *IEEE Transactions on Medical Imaging*, vol. 18, no. 3, pp. 282–290, Mar. 1999.

-
- [8] G. Hamarneh and T. Gustavsson, "Combining snakes and active shape models for segmenting the human left ventricle in echocardiographic images," in *Proceedings of Computers in Cardiology*, vol. 27. Cambridge, USA: IEEE Computer Society Press, September 24-27, 2000, pp. 115–118.
- [9] G. Jacob, J. Noble, C. Behrenbruch, A. Kelion, and A. Banning, "A shape-space-based approach to tracking myocardial borders and quantifying regional left-ventricular function applied in echocardiography," *IEEE Transactions on Medical Imaging*, vol. 21, no. 3, pp. 226–238, Mar. 2002.
- [10] T. Cootes, G. Edwards, and C. Taylor, "Active appearance models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 681–685, June 2001.
- [11] J. Bosch, S. Mitchell, B. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J. Reiber, "Automatic segmentation of echocardiographic sequences by active appearance motion models," *IEEE Transactions on Medical Imaging*, vol. 21, no. 11, pp. 1374–1383, Nov. 2002.
- [12] N. Friedland and D. Adam, "Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing," *IEEE Transactions on Medical Imaging*, vol. 8, no. 4, pp. 344–353, Dec. 1989.
- [13] E. Ashton and K. Parker, "Multiple resolution Bayesian segmentation of ultrasound images," *Ultrasonic Imaging*, vol. 17, no. 4, pp. 291–304, Oct. 1995.
- [14] F. Yeung, S. Levinson, D. Fu, and K. Parker, "Feature-adaptive motion tracking of ultrasound image sequences using a deformable mesh," *IEEE Transactions on Medical Imaging*, vol. 17, no. 6, pp. 945–956, Dec. 1998.
- [15] M. Ledesma-Carbayo, J. Kybic, M. Desco, A. Santos, and M. Unser, "Cardiac motion analysis from ultrasound sequences using non-rigid registration," in *Proceedings of the Fourth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'01)*, ser. Lecture Notes in Computer Science, W. Niessen and M. Viergever, Eds., vol. 2208. Utrecht, The Netherlands: Springer, October 14-17, 2001, pp. 889–896.
- [16] B. Horn and B. Schunk, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.

-
- [17] G. Mailloux, F. Langlois, P. Simard, and M. Bertrand, "Restoration of the velocity field of the heart from two-dimensional echocardiograms," *IEEE Transactions on Medical Imaging*, vol. 8, no. 2, pp. 143–153, June 1989.
- [18] G. Zini, A. Sarti, and C. Lamberti, "Application of continuum theory and multi-grid methods to motion evaluation from 3D echocardiography," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, no. 2, pp. 297–308, Mar. 1997.
- [19] B. Lucas and T. Kanade, "An iterative image restoration technique with an application to stereo vision," *Proc. DARPA IU Workshop*, pp. 121–130, 1981.
- [20] Y. Chunke, K. Terada, and S. Oe, "Motion analysis of echocardiograph using optical flow method," *IEEE International Conference on Systems, Man and Cybernetics*, vol. 1, pp. 672–677, 1996.
- [21] P. Baraldi, A. Sarti, C. Lamberti, A. Prandini, and F. Sgallari, "Evaluation of differential optical flow techniques on synthesized echo images," *IEEE Transactions on Biomedical Engineering*, vol. 43, no. 3, pp. 259–272, Mar. 1996.
- [22] E. Braunwald, D. Zipes, and P. Libby, Eds., *Heart Disease: A Textbook of Cardiovascular Medicine*, 6th ed. W. B. Saunders, 2001, 2-Volume Set.
- [23] E. Braunwald, A. Fauci, D. Kasper, S. Hauser, D. Longo, and J. Jameson, Eds., *Harrison's Principles of Internal Medicine*, 15th ed. McGraw-Hill Professional, 2001.
- [24] N. Pandian, S. Koyanagi, D. Skorton, S. Collins, C. Eastham, M. Marcus, and R. Kerber, "Relations between two-dimensional echocardiographic wall thickening abnormalities, infarct size and coronary risk area in normal and hypertrophied myocardium in dogs," *American Journal of Cardiology*, vol. 52, pp. 1318–1325, 1983.
- [25] G. Kan, C. Visser, J. Koolen, and A. Dunning, "Short and long term predictive value of admission wall motion score in acute myocardial infarction. A cross sectional echocardiographic study of 345 patients," *British Heart Journal*, vol. 56, no. 5, pp. 422–427, Nov. 1986.
- [26] F. Flachskampf, *Kursbuch Echokardiographie*. Thieme, Stuttgart, Germany, 2001, in German.

- [27] T. Faber, M. Akers, R. Peshock, and J. Corbett, "Three-dimensional motion and perfusion quantification in gated single-photon emission computed tomograms," *J. Nucl. Med.*, vol. 32, no. 12, pp. 2311–2317, Dec. 1991.
- [28] G. Germano, J. Erel, H. Kiat, P. Kavanagh, and D. Berman, "Quantitative LVEF and qualitative regional function from gated thallium-201 perfusion SPECT," *J. Nucl. Med.*, vol. 38, no. 5, pp. 749–754, May 1997.
- [29] G. Klein and R. Huesman, "Four-dimensional processing of deformable cardiac PET data," *Med. Im. Anal.*, vol. 6, no. 1, pp. 29–46, Mar. 2002.
- [30] S. Song and R. Leahy, "Computation of 3-D velocity fields from 3-D cine CT images of a human heart," *IEEE Transactions on Medical Imaging*, vol. 10, no. 3, pp. 295–306, Sept. 1991.
- [31] A. Taratorin and S. Sideman, "Constrained detection of left ventricular boundaries from cine CT images of human hearts," *IEEE Transactions on Medical Imaging*, vol. 12, no. 3, pp. 521–533, Sept. 1993.
- [32] M. Kachelriess, S. Ulzheimer, and W. Kalender, "ECG-correlated imaging of the heart with subsecond multislice spiral CT," *IEEE Transactions on Medical Imaging*, vol. 19, no. 9, pp. 888–901, Sept. 2000.
- [33] G. Wright, "Magnetic resonance imaging," *IEEE Signal Processing Magazine*, vol. 14, no. 1, pp. 56–66, Jan. 1997.
- [34] E. Zerhouni, D. Parish, W. Rogers, A. Yang, and E. Shapiro, "Human heart: tagging with MR imaging—a method for noninvasive assessment of myocardial motion," *Radiology*, vol. 169, no. 1, pp. 59–63, Oct. 1988.
- [35] L. Axel and L. Dougherty, "Heart wall motion: improved method of spatial modulation of magnetization for MR imaging," *Radiology*, vol. 172, no. 2, pp. 349–350, Aug. 1989.
- [36] P. Selskog, E. Heiberg, T. Ebbes, L. Wigstrom, and M. Karlsson, "Kinematics of the heart: strain-rate imaging from time-resolved three-dimensional phase contrast MRI," *IEEE Transactions on Medical Imaging*, vol. 21, no. 9, pp. 1105–1109, Sept. 2002.
- [37] A. E. Weyman, *Principles and Practice of Echocardiography*, 2nd ed. Lea & Febiger, 1994.

- [38] M. Marcus, E. Braunwald, M. Marcus, H. Schelbert, D. Skorton, and G. Wolf, Eds., *Marcus Cardiac Imaging: A Companion to Braunwald's Heart Disease*, 2nd ed. W. B. Saunders, 1996, 2-Volume Set.
- [39] J. Quistgaard, "Signal acquisition and processing in medical diagnostic ultrasound," *IEEE Signal Processing Magazine*, vol. 14, no. 1, pp. 67–74, Jan. 1997.
- [40] M. Arigovindan, M. Sühling, P. Hunziker, and M. Unser, "Variational image reconstruction from arbitrarily spaced samples: a fast multiresolution spline solution," *IEEE Transactions on Image Processing*, 2004, in press.
- [41] J. Raichlen, S. Trivedi, G. Herman, M. Sutton, and N. Reichek, "Dynamic three-dimensional reconstruction of the left ventricle from two-dimensional echocardiograms," *J. Am. Coll. Cardiol.*, vol. 8, no. 2, pp. 364–370, Aug. 1986.
- [42] N. Pandian, J. Roelandt, N. Nanda, L. Sugeng, Q. Cao, J. Azevedo, S. Schwartz, M. Vannan, A. Ludomirski, G. Marx, and M. Vogel, "Dynamic three-dimensional echocardiography: methods and clinical potential," *Echocardiography*, vol. 11, no. 3, pp. 237–259, 1994.
- [43] L. Sugeng, L. Weinert, K. Thiele, and R. Lang, "Real-time three-dimensional echocardiography using a novel matrix array transducer," *Echocardiography*, vol. 20, no. 7, pp. 623–635, Oct. 2003.
- [44] J. Doppler, "Über das farbige Licht der Doppelsterne und einiger anderer Gestirne des Himmels," *Abhandlungen der königlichen böhmischen Gesellschaft der Wissenschaften*, vol. 2, pp. 465–482, 1842.
- [45] J. Jensen, *Estimation of Blood Velocities Using Ultrasound: A Signal Processing Approach*. Cambridge University Press, New York, 1996.
- [46] M. Asmi and M. Walsh, *A Practical Guide to Echocardiography*. Chapman & Hall Medical, London, 1995.
- [47] N. Schiller, P. Shah, M. Crawford, A. DeMaria, R. Devereux, H. Feigenbaum, H. Gutgesell, N. Reichek, D. Sahn, I. Schnittger, *et al.*, "Recommendations for quantitation of the left ventricle by two-dimensional echocardiography," *J. Am. Soc. Echocardiogr.*, vol. 2, no. 5, pp. 358–367, Sept.-Oct. 1989.

- [48] M. Sühling, M. Arigovindan, C. Jansen, P. Hunziker, and M. Unser, "Myocardial motion analysis from B-mode echocardiograms," *IEEE Transactions on Image Processing*, 2004, in press.
- [49] P. Hunziker, M. Picard, N. Jander, M. Scherrer-Crosbie, M. Pfisterer, and P. Buser, "Regional wall motion assessment in stress echocardiography by tissue Doppler bull's-eyes." *J. Am. Soc. Echocardiography*, vol. 12, no. 3, pp. 196–202, Mar. 1999.
- [50] M. Sonka and J. Fitzpatrick, Eds., *Handbook of Medical Imaging*. SPIE PRESS, 2000, vol. 2, ch. 12: Cardiac Image Analysis: Motion and Deformation.
- [51] M. Unser, "Splines: A perfect fit for signal and image processing," *IEEE Signal Processing Magazine*, vol. 16, no. 6, pp. 22–38, Nov. 1999.
- [52] X. Zong, A. Laine, and E. Geiser, "Speckle reduction and contrast enhancement of echocardiograms via multiscale nonlinear processing," *IEEE Transactions on Medical Imaging*, vol. 17, no. 4, pp. 532–540, Aug. 1998.
- [53] S. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, July 1989.
- [54] C. Shannon, "Communication in the presence of noise," *Proceedings of the Institute of Radio Engineers*, vol. 37, no. 1, pp. 10–21, Jan. 1949.
- [55] M. Unser, A. Aldroubi, and M. Eden, "The L_2 -polynomial spline pyramid," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 4, pp. 364–379, Apr. 1993.
- [56] —, "B-Spline signal processing: Part I—Theory," *IEEE Transactions on Signal Processing*, vol. 41, no. 2, pp. 821–833, Feb. 1993.
- [57] —, "B-Spline signal processing: Part II—Efficient design and applications," *IEEE Transactions on Signal Processing*, vol. 41, no. 2, pp. 834–848, Feb. 1993.
- [58] M. Jacob, T. Blu, and M. Unser, "An exact method for computing the area moments of wavelet and spline curves," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 633–642, June 2001.

- [59] J. C. Bamber and R. J. Dickinson, "Ultrasonic B-scanning: a computer simulation," *Phys. Med. Biol.*, vol. 25, pp. 463–479, 1980.
- [60] M. Bertrand and J. Meunier, "Ultrasonic texture motion analysis: theory and simulation," *IEEE Transactions on Medical Imaging*, vol. 14, no. 2, pp. 293–300, June 1995.
- [61] J. Barron, D. Fleet, S. Beauchemin, and T. Burkitt, "Performance of optical flow techniques," *Intl. J. Computer Vision*, vol. 12, pp. 43–77, 1994.
- [62] S. Van Huffel and J. Vandewalle, *The Total Least Squares Problem: Computational Aspects and Analysis*, ser. Frontiers in Applied Mathematics. SIAM, Philadelphia, 1991, vol. 9.
- [63] G. Golub and C. Van Loan, "An analysis of the total least squares problem," *SIAM Journal on Numerical Analysis*, vol. 17, no. 6, pp. 883–893, Dec. 1980.
- [64] M. Sonka and J. Fitzpatrick, Eds., *Handbook of Medical Imaging*. SPIE PRESS, 2000, vol. 2, ch. 11: Echocardiography.
- [65] M. Mühlich and R. Mester, "The role of total least squares in motion analysis," in *Proceedings of the 5th European Conference on Computer Vision (ECCV)*, ser. Lecture Notes in Computer Science, H. Burkhardt and B. Neumann, Eds., vol. 1407. Freiburg, Germany: Springer, June 2-6, 1998, pp. 305–321.
- [66] B. Jähne, *Digitale Bildverarbeitung*. Springer, Berlin, 1989, in German.
- [67] J. Bigün, G. Granlund, and J. Wiklund, "Multidimensional orientation estimation with applications to texture analysis and optical flow," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 8, pp. 775–790, Aug. 1991.
- [68] A. Sayed, V. Nascimento, and S. Chandrasekaran, "Estimation and control with bounded data uncertainties," *Linear Algebra and Its Applications*, vol. 284, pp. 259–306, Nov. 1998.
- [69] S. Chandrasekaran, G. Golub, M. Gu, and A. Sayed, "An efficient algorithm for a bounded errors-in-variables model," *SIAM J. Mat. Anal. Appl.*, vol. 20, no. 4, pp. 839–859, Oct. 1999.

- [70] C. Paige and Z. Strakoš, "Scaled total least squares fundamentals," *Numerische Mathematik*, vol. 91, no. 1, pp. 117–146, Mar. 2002.
- [71] M. Mühlich and R. Mester, "A considerable improvement in non-iterative homography estimation using TLS and equilibration," *Pattern Recognition Letters*, vol. 22, no. 11, pp. 1181–1189, Sept. 2001.
- [72] M. Sühling, M. Arigovindan, C. Jansen, P. Hunziker, and M. Unser, "Bi-modal myocardial motion analysis from B-mode and tissue Doppler ultrasound," in *Proceedings of the Second 2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'04)*, Arlington VA, USA, April 15-18, 2004, pp. 308–311.
- [73] W. McDicken, G. Sutherland, C. Moran, and L. Gordon, "Colour doppler velocity imaging of the myocardium," *Ultrasound Med. Biol.*, vol. 18, no. 6-7, pp. 651–654, 1992.
- [74] G. Sutherland, M. Steward, K. Grounstroem, C. Moran, A. Fleming, F. Guell-Peris, R. Riemersma, L. Fenn, K. Fox, and W. Mc Dicken, "Color doppler myocardial imaging: a new technique for assessment of myocardial function," *J. Am. Soc. Echocardiogr.*, vol. 7, no. 5, pp. 441–458, Sept.–Oct. 1994.
- [75] K. Miyatake, M. Yamagishi, N. Tanaka, M. Uematsu, N. Yamazaki, Y. Mine, A. Sano, and M. Hirama, "New method of evaluating left ventricular wall motion by color-coded tissue Doppler imaging: in vitro and in vivo studies," *J. Am. Coll. Cardiol.*, vol. 25, no. 3, pp. 717–724, Mar. 1995.
- [76] J. Gorcsan, V. Gulati, W. Mandarino, and W. Katz, "Color-coded measures of myocardial velocity throughout the cardiac cycle by tissue Doppler imaging to quantify regional left ventricular function," *Am. Heart J.*, vol. 131, no. 6, pp. 1203–1213, June 1996.
- [77] L. Galiuto, G. Ignone, and A. DeMaria, "Contraction and relaxation velocities of the normal left ventricle using pulsed-wave tissue Doppler echocardiography," *Am. J. Cardiology*, vol. 81, no. 5, pp. 609–614, Mar. 1998.
- [78] W. Katz, V. Gulati, C. Mahler, and J. Gorcsan, "Quantitative evaluation of the segmental left ventricular response to dobutamine stress by tissue Doppler echocardiography," *Am. J. Cardiology*, vol. 79, no. 8, pp. 1036–1042, Apr. 1997.

- [79] G. Derumeaux, M. Ovize, J. Loufoua, X. Andre-Fouet, Y. Minaire, A. Cribier, and B. Letac, "Doppler tissue imaging quantitates regional wall motion during myocardial ischemia and reperfusion," *Circulation*, vol. 97, no. 19, pp. 1970–1977, May 1998.
- [80] J. D'hooge, A. Heimdal, F. Jamal, T. Kukulski, B. Bijnens, F. Rademakers, L. Hatle, P. Suetens, and G. R. Sutherland, "Regional strain and strain rate measurements by cardiac ultrasound: Principles, implementation and limitations," *European Journal of Echocardiography*, vol. 1, no. 3, pp. 154–170, Sept. 2000.
- [81] A. Fleming, X. Xia, W. McDicken, G. Sutherland, and L. Fenn, "Myocardial velocity gradients detected by doppler imaging system," *British J. Radiology*, vol. 67, pp. 679–688, July 1994.
- [82] M. Uematsu, K. Miyatake, N. Tanaka, H. Matsuda, A. Sano, N. Yamazaki, M. Hirama, and M. Yamagishi, "Myocardial velocity gradient as a new indicator of regional left ventricular contraction: detection by a two-dimensional tissue doppler imaging technique," *J. Am. Coll. Cardiology*, vol. 26, no. 1, pp. 217–223, July 1995.
- [83] B. Guth, R. Savage, F. White, A. Hagan, L. Samtoy, and C. Bloor, "Detection of ischemic wall dysfunction: Comparison between M-mode echocardiography and sonomicrometry," *Am. Heart J.*, vol. 107, pp. 449–457, 1984.
- [84] H. Kanai, H. Hasegawa, N. Chubachi, Y. Koiwa, and M. Tanaka, "Noninvasive evaluation of local myocardial thickening and its color-coded imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, no. 4, pp. 752–768, July 1997.
- [85] M. Feistauer, *Mathematical Methods in Fluid Dynamics*, ser. Pitman Monographs and Surveys in Pure and Applied Mathematics. Longman Scientific & Technical, Harlow, 1993, vol. 67.
- [86] T. Denney, Jr. and J. Prince, "Reconstruction of 3-D left ventricular motion from planar tagged cardiac MR images: an estimation theoretic approach," *IEEE Transactions on Medical Imaging*, vol. 14, no. 4, pp. 625–635, Dec. 1995.
- [87] M. Kowalski, T. Kukulski, F. Jamal, J. D'hooge, F. Weidemann, F. Rademakers, B. Bijnens, L. Hatle, and G. Sutherland, "Can natural strain and

- strain rate quantify regional myocardial deformation? A study in healthy subjects,” *Ultrasound in Medicine and Biology*, vol. 27, no. 8, pp. 1087–1097, Aug. 2001.
- [88] J. Bland and D. Altman, “Statistical methods for assessing agreement between two methods of clinical measurement,” *Lancet*, vol. 1, no. 8476, pp. 307–310, 1986.
- [89] W. Rosenkrantz, *Introduction to Probability and Statistics for Scientists and Engineers*. McGraw-Hill Publishing, New York, 1997.
- [90] Y. Hochberg and A. Tamhane, *Multiple Comparison Procedures*. Wiley, 1987.
- [91] M. Hagan, H. Demuth, and M. Beale, *Neural Network Design*. Boston, MA: PWS Publishing, 1996.
- [92] J. Cohen, “A coefficient of agreement for nominal scales,” *Educational and Psychological Measurement*, vol. 20, pp. 27–46, 1960.
- [93] J. Fleiss, *Statistical Methods for Rates and Proportions*. John Wiley and Sons, New York, 1981.
- [94] C. A. Gibbons Kroeker, J. V. Tyberg, and R. Beyar, “Effects of Ischemia on Left Ventricular Apex Rotation: An Experimental Study in Anesthetized Dogs,” *Circulation*, vol. 92, no. 12, pp. 3539–3548, 1995.
- [95] M. Sühling, M. Arigovindan, P. Hunziker, and M. Unser, “Multiresolution moment filters: Theory and applications,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 484–495, Apr. 2004.
- [96] M.-K. Hu, “Visual pattern recognition by moment invariants,” *IRE Trans. Inform. Theory*, pp. 179–187, Feb. 1962.
- [97] C. Teh and R. Chin, “On image analysis by the methods of moments,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 4, pp. 496–513, July 1988.
- [98] S. Dudani, K. Breeding, and R. McGhee, “Aircraft identification by moment invariants,” *IEEE Transactions on Computers*, vol. 26, no. 1, pp. 39–46, Jan. 1977.

-
- [99] Y. Abu-Mostafa and D. Psaltis, "Recognitive aspects of moment invariants," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 698–706, Nov. 1984.
- [100] M. Teague, "Image analysis via the general theory of moments," *J. Optical Soc. Am.*, vol. 70, pp. 920–930, Aug. 1980.
- [101] S. Liao and M. Pawlak, "On image analysis by moments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 3, pp. 254–266, Mar. 1996.
- [102] A. Khotanzad and Y. Hong, "Invariant image recognition by Zernike moments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 5, pp. 489–497, May 1990.
- [103] A. Bhatia and E. Wolf, "On the circle polynomials of Zernike and related orthogonal sets," *Proc. Cambridge Philosophical Society*, vol. 50, pp. 40–48, 1954.
- [104] T. Reiss, *Recognizing planar objects using invariant image features. Lecture notes in computer science.* Springer, Berlin, 1993, vol. 676.
- [105] B. Li and J. Shen, "Fast calculation of local moments and application to range segmentation," *Proceedings International Conference on Pattern Recognition, The Haage*, vol. C, pp. 298–301, 1992.
- [106] M. Tuceryan, "Moment-based texture segmentation," *Pattern Recognition Letters*, vol. 15, pp. 659–668, July 1994.
- [107] Q. Gao and F. Yin, "Two-dimensional direction-based interpolation with local centered moments," *Graphical Models and Image Processing*, vol. 61, no. 6, pp. 323–339, Nov. 1999.
- [108] J. Martinez and F. Thomas, "Efficient computation of local geometric moments," *IEEE Transactions on Image Processing*, vol. 11, no. 9, pp. 1102–1111, 2002.
- [109] J.-W. Lin, A. Laine, and S. Bergmann, "Improving PET-based physiological quantification through methods of wavelet denoising," *IEEE Transactions on Biomedical Engineering*, vol. 48, no. 2, pp. 202–212, Feb. 2001.

- [110] R. Nowak, "Wavelet-based Rician noise removal for magnetic resonance imaging," *IEEE Transactions on Image Processing*, vol. 8, no. 10, pp. 1408–1419, Oct. 1999.
- [111] A. Laine, S. Schuler, J. Fan, and W. Huda, "Mammographic feature enhancement by multiscale analysis," *IEEE Transactions on Medical Imaging*, vol. 13, no. 4, pp. 725–740, Dec. 1994.
- [112] J.-P. Antoine, D. Barache, R. Cesar Jr., and L. da F. Costa, "Multiscale shape analysis using the continuous wavelet transform," in *Proceedings of the 1996 IEEE International Conference on Image Processing (ICIP'96)*, vol. I, Lausanne, Switzerland, Sept. 16–19 1996, pp. 291–294.
- [113] A. Savitzky and M. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.
- [114] M. Shensa, "The discrete wavelet transform: wedding the à trous and Mallat algorithms," *IEEE Transactions on Signal Processing*, vol. 40, no. 10, pp. 2464–2482, Oct. 1992.
- [115] M. Unser and T. Blu, "Wavelet theory demystified," *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 470–483, Feb. 2003.
- [116] J. Dubochet, M. Adrian, J. Chang, J. Homo, J. Lepault, A. McDowell, and P. Schultz, "Cryo-electron microscopy of vitrified specimens," *Quarterly Review of Biophysics*, vol. 21, no. 2, pp. 129–228, 1988.
- [117] T. Blu and M. Unser, "The fractional spline wavelet transform: Definition and implementation," in *Proceedings of the Twenty-Fifth IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'00)*, vol. I, Istanbul, Turkey, June 5–9 2000, pp. 512–515.
- [118] J. S. Lim, *Two-Dimensional Signal and Image Processing*. Prentice Hall, 1990, pp. 536–540.

Curriculum Vitæ

Michael Sühling
Biomedical Imaging Group
STI/IOA/LIB, Bldg. BM 4.136
Swiss Federal Institute of Technology Lausanne
CH-1015 Lausanne
Switzerland

Tel.: +41 21 693 51 35
URL: <http://bigwww.epfl.ch/suehling/>
E-mail: michael.suehling@epfl.ch



PERSONAL DETAILS

Date of birth: October 18, 1974
Place of birth: Bocholt, Germany
Nationality: German

EDUCATION

Sept. 2000 to June 2004 PhD-student in biomedical image analysis, Swiss Federal Institute of Technology Lausanne (EPFL), Lausanne, Switzerland:
• PhD-dissertation: “*Myocardial Motion and Deformation Analysis from Echocardiograms*”, Advisor: Prof. M. Unser

- Oct. 1994 to July 2000 Diploma in Technomathematics, University of Kaiserslautern, Germany:
- Participation in the European study program “Mathematics for Industry”
 - Study period at the Chalmers University of Technology, Gothenburg, Sweden
- June 1994 Abitur (diploma from German secondary school qualifying for university admission), Gymnasium Mariengarden, Borken; awarded by best-of-class prize

PRACTICAL EXPERIENCE

- Sept. 2000 to June 2004 Research assistant at the Biomedical Imaging Group, EPFL:
- Supervision of several student semester projects in image processing
 - Teaching assistant for:
 - Image Processing 1&2
 - Signals and Systems 1&2
 - Reviewer for:
 - IEEE Transactions on Image Processing
 - IEEE Transactions on Medical Imaging
- March 1999 to May 2000 Part-time research assistant at the Fraunhofer Institute for Industrial Mathematics (ITWM), Kaiserslautern, Germany
- Aug. 1998 to Oct. 1998 Internship at Summus Limited, Columbia-SC, USA

SPECIAL SKILLS

- Languages:
- German (native speaker)
 - English (fluently spoken and written)
 - French (working knowledge)
 - Swedish (basic knowledge)

Computer skills: • Programming: C/C++, Java, Matlab, Khoros, Oberon, Object Modeling Technique (OMT)
 • Operating systems: Windows, Mac OS, Linux
 • Publishing: T_EX, HTML, Acrobat, Office

AWARDS

May 2003 “Menarini Prize for Echocardiography”, Annual Meeting of the Swiss Society of Cardiology

Dec. 2002 “Young Investigators’ Award” together with C. Jansen, MD, Euroecho 6 Conference, Munich, Germany

INTERESTS

• Jogging, playing soccer, skiing, playing keyboard
 • Member of student organization IAESTE Lausanne (International Association for the Exchange of Students for Technical Experience)

Publications

Journal Papers

- [1] M. Arigovindan, M. Sühling, P. Hunziker, and M. Unser, “Variational image reconstruction from arbitrarily spaced samples: a fast multiresolution spline solution,” *IEEE Transactions on Image Processing*, 2004, in press.
- [2] M. Sühling, M. Arigovindan, C. Jansen, P. Hunziker, and M. Unser, “Myocardial motion analysis from B-mode echocardiograms,” *IEEE Transactions on Image Processing*, 2004, in press.
- [3] M. Sühling, M. Arigovindan, P. Hunziker, and M. Unser, “Multiresolution moment filters: Theory and applications,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 484–495, Apr. 2004.

Journal Papers in Review

- [1] M. Sühling, C. Jansen, M. Arigovindan, P. Buser, S. Marsch, M. Unser, and P. Hunziker, “Multiscale Motion Mapping—a novel computer vision

technique for quantitative, objective echocardiographic motion measurement independent of Doppler: First clinical description and validation,” *Circulation*, 2004, accepted with minor revision.

- [2] M. Ledesma-Carbayo, J. Kybic, M. Desco, A. Santos, M. Sühling, P. Hunziker, and M. Unser, “Spatio-temporal non-rigid registration for ultrasound cardiac motion detection,” *IEEE Transactions on Medical Imaging*, 2004, in review.

Conference Papers

- [1] M. Sühling, M. Arigovindan, C. Jansen, P. Hunziker, and M. Unser, “Bimodal myocardial motion analysis from B-mode and tissue Doppler ultrasound,” in *Proceedings of the Second 2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'04)*, Arlington VA, USA, April 15-18, 2004, pp. 308–311.
- [2] C. Jansen, M. Arigovindan, M. Sühling, S. Marsch, M. Unser, and P. Hunziker, “Multidimensional, multistage wavelet footprints: A new tool for image segmentation and feature extraction in medical ultrasound,” in *Progress in Biomedical Optics and Imaging, vol. 4, no. 23*, ser. Proceedings of the SPIE International Symposium on Medical Imaging: Image Processing (MI'03), M. Sonka and J. Fitzpatrick, Eds., vol. 5032, San Diego CA, USA, February 17-20, 2003, pp. 762–767, part II.
- [3] M. Sühling, M. Arigovindan, C. Jansen, P. Hunziker, and M. Unser, “Myocardial motion analysis and visualization from echocardiograms,” in *Progress in Biomedical Optics and Imaging, vol. 4, no. 23*, ser. Proceedings of the SPIE International Symposium on Medical Imaging: Image Processing (MI'03), M. Sonka and J. Fitzpatrick, Eds., vol. 5032, San Diego CA, USA, February 17-20, 2003, pp. 306–313, part I.
- [4] M. Arigovindan, M. Sühling, P. Hunziker, and M. Unser, “Multigrid image reconstruction from arbitrarily spaced samples,” in *Proceedings of the 2002 IEEE International Conference on Image Processing (ICIP'02)*, vol. III, Rochester NY, USA, September 22-25, 2002, pp. 381–384.
- [5] M. Sühling, M. Arigovindan, P. Hunziker, and M. Unser, “Multiresolution moment filters,” in *Proceedings of the 2002 IEEE International Conference on Image Processing (ICIP'02)*, vol. I, Rochester NY, USA, Sept. 22–25 2002, pp. 393–396.

- [6] M. Ledesma-Carbayo, J. Kybic, M. Sühling, P. Hunziker, M. Desco, A. Santos, and M. Unser, “Cardiac ultrasound motion detection by elastic registration exploiting temporal coherence,” in *Proceedings of the First 2002 IEEE International Symposium on Biomedical Imaging: Macro to Nano (ISBI’02)*, vol. II, Washington DC, USA, July 7-10, 2002, pp. 585–588.
- [7] M. Sühling, M. Arigovindan, P. Hunziker, and M. Unser, “Motion analysis of echocardiograms using a local-affine, spatio-temporal model,” in *Proceedings of the First 2002 IEEE International Symposium on Biomedical Imaging: Macro to Nano (ISBI’02)*, vol. II, Washington DC, USA, July 7–10 2002, pp. 573–576.

Conference Abstracts

- [1] M. Sühling, C. Jansen, S. Aschkenasy, P. Buser, S. Marsch, M. Unser, and P. Hunziker, “Vector motion analysis for objective, quantitative analysis of echocardiograms: First large-scale clinical validation in 100 patients,” in *Actes de l’assemblée annuelle commune, Société Suisse de Cardiologie, Société Suisse d’Angiologie*, Basel, Switzerland, June 2-4, 2004, vol. 7, Supplementum 6, pp. 26 S.
- [2] M. Sühling, M. Unser, C. Jansen, M. Arigovindan, P. Buser, S. Marsch, and P. Hunziker, “Exploring the complexity of cardiac motion by multi-scale motion mapping, a novel quantitative echocardiographic technique,” in *Twenty-Fifth Annual Congress of the European Society of Cardiology (ESC’03)*, ser. European Heart Journal, vol. 24, abstr. supp. 1, August/September 2003, Vienna, Austria, August 30-September 2, 2003, p. 376.
- [3] P. Hunziker, M. Sühling, M. Arigovindan, C. Jansen, A. Linka, P. Buser, S. Marsch, and M. Unser, “Exploring the complexity of cardiac motion by multiscale motion mapping,” in *Actes de l’assemblée annuelle commune, Société Suisse de Cardiologie, Société Suisse de Chirurgie Thoracique et Cardio-Vasculaire, Société Suisse de Médecine Intensive*, vol. 6, supp. 5, Lausanne VD, Switzerland, May 8-10, 2003, p. 17 S.
- [4] M. Sühling, M. Arigovindan, C. Jansen, A. Linka, P. Buser, C. Heidegger, S. Marsch, P. Hunziker, and M. Unser, “Direct measurement of myocardial thickening by means of computer vision,” in *Actes de l’assemblée annuelle*

- commune, Société Suisse de Cardiologie, Société Suisse de Chirurgie Thoracique et Cardio-Vasculaire, Société Suisse de Médecine Intensive*, vol. 6, Supplementum 5, Lausanne VD, Switzerland, May 8-10, 2003, p. 16 S.
- [5] C. Jansen, M. Arigovindan, M. Sühling, M. Unser, S. Marsch, and P. Hunziker, “Computer vision in echocardiography: A new technology for autonomous segmentation of cardiac structures,” in *Fifty-Second Annual Scientific Session of the American College of Cardiology (ACC’03)*, ser. Journal of the American College of Cardiology, vol. 41, no. 6, supp. A, March 2003, Chicago IL, USA, March 30-April 2, 2003, p. 538A.
- [6] C. Jansen, M. Sühling, M. Arigovindan, M. Unser, P. Buser, and P. Hunziker, “Multiscale motion mapping (“triple-M imaging”): A novel echo imaging mode featuring quantitative echo motion analysis,” in *Fifty-Second Annual Scientific Session of the American College of Cardiology (ACC’03)*, ser. Journal of the American College of Cardiology, vol. 41, no. 6, supp. A, March 2003, Chicago IL, USA, March 30-April 2, 2003, p. 455A.
- [7] C. Jansen, L. Zeng, M. Arigovindan, M. Sühling, S. Marsch, and P. Hunziker, “Wavelet- versus JPEG- compression of echocardiograms,” in *Fifty-Second Annual Scientific Session of the American College of Cardiology (ACC’03)*, ser. Journal of the American College of Cardiology, vol. 41, no. 6, supp. A, March 2003, Chicago IL, USA, March 30-April 2, 2003, p. 449A.
- [8] C. Jansen, M. Sühling, M. Arigovindan, M. Unser, S. Marsch, and P. Hunziker, “Multiscale motion mapping (triple-M imaging) for colour-coded analysis of stress echocardiograms,” in *Fifty-Second Annual Scientific Session of the American College of Cardiology (ACC’03)*, ser. Journal of the American College of Cardiology, vol. 41, no. 6, supp. A, March 2003, Chicago IL, USA, March 30-April 2, 2003, p. 426A.
- [9] C. Jansen, M. Sühling, M. Arigovindan, P. Buser, M. Unser, and P. Hunziker, “Multiscale motion mapping (“triple-M imaging”): A novel echo imaging mode featuring quantitative echo motion display independent from borders and Doppler,” in *Sixth Annual and Plenary Meeting of the Working Group on Echocardiography of the European Society of Cardiology (EUROECHO6’02)*, ser. European Journal of Echocardiography, vol. 3, supp. 1, Munich, Germany, December 4-7, 2002, p. 122.

-
- [10] T. Blu, M. Sühling, P. Thévenaz, and M. Unser, “Approximation order: Why the asymptotic constant matters,” in *Second Pacific Rim Conference on Mathematics (PRCM’01)*, Taipei, Taiwan ROC, January 4-8, 2001, pp. II.3–II.4.

