



Universiteit
Leiden
The Netherlands

Automated analysis of 3D echocardiography

Stralen, M. van

Citation

Stralen, M. van. (2009, February 25). *Automated analysis of 3D echocardiography*. *ASCI dissertation series*. Retrieved from <https://hdl.handle.net/1887/13521>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/13521>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/13521> holds various files of this Leiden University dissertation.

Author: Stralen, M. van

Title: Automated analysis of 3D echocardiography

Issue date: 2009-02-25

Automated analysis of 3D echocardiography

10.1016/
j.ultrasmedbio.20
07.03.007

Marijn van Stralen

| Colophon

Automated analysis of 3D echocardiography
Stralen, Marijn van

ISBN: 978-90-8559-493-2

Printed by Optima Grafische Communicatie, Rotterdam, the Netherlands

© 2009 M. van Stralen, Leiden, the Netherlands

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the copyright owner.

Automated analysis of 3D echocardiography

Proefschrift


ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van Rector Magnificus prof.mr. P.F. van der Heijden,
volgens besluit van het College voor Promoties
te verdedigen op woensdag 25 februari 2009
klokke 15:00 uur

door

Marijn van Stralen
geboren te Roermond
in 1980

Promotiecommissie

Promotores:	Prof. dr. ir. J.H.C. Reiber Prof. dr. ir. A.F.W. van der Steen (Erasmus MC Rotterdam)
Copromotor:	Dr. ir. J.G. Bosch (Erasmus MC Rotterdam)
Referent:	Prof. dr. W.J. Niessen (Erasmus MC Rotterdam)
Overige leden:	Prof. dr. ir. N. de Jong (Erasmus MC Rotterdam) Prof. dr. M.J. Schalij



This work was carried out in the ASCI graduate school.
ASCI dissertation series number 171

This study was part of ICIN project number 47. It was financially supported by the Technology Program of the Dutch Ministry of Economic Affairs (SenterNovem IOP, grant IBVC-02003). Chapter 6 was also supported by the Dutch Technology Foundation STW (grant 06666), applied science division of NWO.

Financial support by the Netherlands Heart Foundation for the publication of this thesis is gratefully acknowledged.

Financial support for the publication of this thesis was kindly provided by:

- ★ Lecoecur Electronique
- ★ Medis medical imaging systems bv
- ★ Oldelft Ultrasound
- ★ TomTec Imaging Systems GmbH

Contents

1	Introduction	1
	Motivation, 2 3D echocardiography, 3 Digital image analysis, 12 Outline of this thesis, 16.	
2	Semi-automatic endocardial border detection for left ventricular volume estimation in 3D echocardiography	19
	Introduction, 20 Methods, 23 Results, 30 Discussion, 32 Conclusions and future work, 35.	
3	Automated tracking of the mitral valve annulus motion in apical echocardiographic images using multidimensional dynamic programming	37
	Introduction, 38 Materials and methods, 39 Results, 47 Discussion, 53 Conclusions, 56.	
4	Interpolation of irregularly distributed sparse 4D ultrasound data using normalized convolution	57
	Introduction, 58 Methods, 62 Experiments, 66 Results, 69 Discussion, 70 Conclusions, 71.	
5	Automatic time continuous detection of the left ventricular long axis and the mitral valve plane in 3D echocardiography	73
	Introduction and literature, 74 Materials and methods, 75 Results, 84 Discussion, 91 Conclusions, 95.	

6 Automated left ventricular volume estimation in 3D echocardiography using active appearance models 97
Introduction, 98 | Active appearance models, 100 | Methods, 104 | Experiments and results, 112 | Discussion, 119 | Conclusions, 128.

7 Discussion and conclusions 131
Research objective, 132 | Contributions, 132 | Discussion, 133 | Recommendations for future work, 135 | Conclusions, 136.

8 Summary 139

9 Samenvatting 143

Bibliography 147

Publications 159

Dankwoord 165

Curriculum Vitae 167

Introduction

1

1.1 | Motivation

cardiovascular disease Cardiovascular disease has been the number one cause of death in the world for the last decades and is projected to remain the leading cause of death [WHO 2007]. Cardiovascular disease encompasses disorders of the heart and blood vessels, either originated at birth (congenital heart disease) or developed during life. Among the latter are atherosclerosis, such as coronary artery disease (possibly causing a heart attack) and cerebrovascular disease (causing stroke), arrhythmias, hypertension, heart failure and many other diseases. Major (modifiable) risk factors for cardiovascular disease include unhealthy diet, physical inactivity and tobacco use.

left ventricle of the heart In this thesis we will focus on the assessment of global functioning of the left ventricle of the heart. This ventricle is responsible for pumping the blood, coming from the lungs, where it is saturated with oxygen, through the whole body. Measurement of the left ventricular volume and function is therefore very important in clinical decision-making, assessment of therapeutic effects and determination of prognosis.

Malfunctioning of the left ventricle may be caused by coronary artery disease, hypertension or arrhythmias. Ischemia, which eventually results in heart failure may cause a wide variety of symptoms. Since in mild cases of heart failure symptoms may be faint and a universally agreed definition is lacking, the disease is often undiagnosed. This may have severe consequences, including even death.

3-dimensional echocardiography While prevention aiming at reduction of the main modifiable risk factors can reduce the number of deaths caused by cardiovascular disease, also a wider availability of diagnostic techniques might improve the treatment of cardiovascular disease. Since the beginning of this century, 3-dimensional echocardiography (3DE) has become available and is getting more widespread across medical centers. 3DE offers a non-invasive, relatively cheap and therefore possibly widely available way to visualize the left ventricle in 3D (fig. 1.1) and to analyze its function. However, manual analysis of these images for quantitative assessment of functional parameters is cumbersome. Therefore, automation of the assessment of left ventricular function is an important step in improving the diagnosis and treatment of cardiovascular disease and reducing its costs.

automating the analysis




Figure 1.1: A 3D ultrasound image of the left ventricle with the endocardial surface in blue. *left*) A 3D rendering of the image with the endocardial surface semi-transparent in red. *top right*) A 2-chamber view of the same data set, with the delineation of the endocardial border. *bottom right*) A short-axis view of the same data set

3D echocardiography | 1.2

Imaging of the heart poses many challenges on imaging modalities. To assess the functional parameters of the heart, the geometry and dynamics of the heart should be imaged in great detail. Therefore, ideally, the full cardiac cycle is imaged in real-time, distinguishing different types of tissue with high spatial and temporal resolution, and with minimal discomfort for the patient, at low costs.

In the past decade, echocardiography (ultrasound imaging of the heart) has been conquering many technological challenges to achieve this goal. It has been developed into a very competitive imaging technique with its own strengths and limitations. In this section we will discuss these characteristics from a technical and clinical point of view.




Figure 1.2: Different acquisition modes for echocardiography. *a)* An M-mode image, showing one acquisition line over time (horizontal axis) *b)* A B-mode image, showing a cross-section of the left ventricle. *c)* A 3D rendering of the left ventricle, acquired by mechanically rotating a phased array transducer around its central image axis

1.2.1 | Ultrasound imaging

principle

Imaging using ultrasound is done by transmitting beams of high frequency sound and recording the resulting echos. Echos come from transitions between materials of different densities, for example air and bone, or blood and tissue. The larger the density difference, the higher the intensity of the echo. Given the speed of sound in the imaged object and the recorded time between sending the sound pulse and receiving the echo we can locate the reflector or scatterer. When imaging tissue, the echos are a result of scattering of the ultrasound beam due to the inhomogeneous nature of the tissue, generating speckle.

Ultrasound is transmitted by applying an electrical field on piezoelectric material (the transducer element), making it vibrate at high frequency and transmitting ultrasound. Receiving echos is basically the same process in reverse. During traversal of the ultrasound beam through the medium, multiple echos from various depths can be recorded, generating an image line (A-mode image). Imaging this line over time gives an M-mode (motion) image (fig. 1.2a).

making an image

A 2-dimensional (2D, B-mode) image is built up from image lines that are recorded by sending and receiving focused sound beams under different angles. This can be done by sweeping the ultrasound beam mechanically (fig. 1.2b). Transmitting a beam under a certain angle can also be done by using multiple sound sources (the transducer elements) and activating them with a time delay between neighboring elements such that a sound wave in the desired direction is created that converges at a certain depth (fig. 1.3). Such an array of transducer elements used for 2D imaging is called a phased-array transducer. Phased-array transducers are most commonly used in 2D echocardiography.




Figure 1.3: Electronic beam steering using a phased array transducer. The transducer elements (dark gray) are activated with such time delays (Δt) that a wave front (solid redline) under the desired angle is created

A few of the most important parameters that determine the quality and resolution of the recorded images are discussed in this paragraph. The material and thickness of the transducer elements determine the resonance frequency and bandwidth of the transducer. The higher the frequency, the smaller the penetration depth, but the higher the axial imaging resolution. The speed of sound is assumed to be constant in human tissue and limits the number of beams that can be sent and received sequentially per time unit, restricting the number of frames that can be imaged per second (the frame rate). The width of the array of transducer elements influences the width of the focussed beam and is therefore a limiting factor in the lateral resolution of the image. In adult 2D echocardiography, typical resolutions are 0.3 mm axially and about 1° laterally.

*major
parameters*

Since a decade, tissue harmonic imaging [Spencer et al. 1998; Tranquart et al. 1999] has been widely adopted in medical ultrasound imaging. Due to nonlinear propagation of the ultrasound wave through the tissue, higher frequencies, harmonic modes of the transmitted signal, are generated. To exploit this phenomenon,

*harmonic
imaging*

transducers need to be designed that are sensitive to these higher frequencies. Previously, transducers were optimized to transmit and receive in the same frequency. Because of the higher frequency, a higher resolution can be achieved in the image [Ward et al. 1997]. Furthermore, harmonic imaging shows some advantages over fundamental imaging which result in clearer images and reduced near-field clutter [Duck 2002; Thomas and Rubin 1998].

1.2.2 | Developments in 3D echocardiography

Conventional 2D echocardiography (B-mode imaging of the heart) allows visualization of a slice of the heart over time and is widely used for assessment of cardiac function. 2D echocardiography allows measurement of left ventricular volume and derived parameters such as ejection fraction, stroke volume and cardiac output. However, assumptions about the left ventricular geometry and the position of the imaged planes in 3D, need to be made.

Therefore, ever since the existence of 2D echocardiography, people have been searching for an extension to 3D, to overcome the limitations of 2D echocardiography [Bruining et al. 2000].

1.2.2.1 | Freehand 3D imaging

First attempts towards 3D echocardiography were made by freehand scanning using a conventional 2D transducer that was registered in 3D either acoustically (so-called spark gap location [Moritz and Shreve 1976]), using a mechanical arm [Dekker et al. 1974] or using electromagnetic spatial locators [Barratt et al. 2001; Raab et al. 1979]. 3D image reconstruction is performed offline by dedicated software. Only the electromagnetic tracking systems eventually made it into the clinic, because of practical limitations of the acoustic and mechanic positioning systems. But still, also electromagnetic freehand 3DE has its limitations. The positioning accuracy is limited and the acquisition is time-consuming and cumbersome. It suffers from motion artifacts as a result of patient movement and breathing. In spite of these restrictions, freehand 3DE has been used till recently [Mannaerts et al. 2003; Varandas et al. 2004] because of its cost-effectiveness.

1.2.2.2 | Mechanical 3D imaging

To shorten acquisition times and improve on irregular coverage of the 3D space by freehand acquisitions, the acquisitions have been automated in several ways.

The first approach is a linear scan of the target space resulting in parallel 2D *linear* images constituting the 3D volume (fig. 1.4a). Pandian et al. explored, among




Figure 1.4: Mechanical 3D scanning modes. *a)* Sweep mode. *b)* Fan-like mode. *c)* Rotational mode

other configurations, the possibilities of computer-controlled serial 2D cardiac tomographic images extensively [Schwartz et al. 1994].

Secondly, mechanical fan-like sweeping of the phased-array has been proposed (fig. 1.4b). In this way, a pyramidal volume can be scanned by moving the transducer in a fan-like arc at prescribed angles [Delabays et al. 1995]. In contrast to the linear scan approach, fan-like movement of the transducer is more suitable for transthoracic echocardiography, because of the limited echo window.

fan-like

Thirdly, stepwise rotational scanning of the volume has been applied (fig. 1.4c), where the phased array is rotated around its central image axis, such that co-axial images are acquired resulting in a conical 3D data set [Pandian et al. 1994]. The number of co-axial planes can be varied to prioritize between the speed of the acquisition and the accuracy of the imaged volume, and the derived clinical parameters [Nosir et al. 1996; Papavassiliou et al. 1998].

stepwise rotating

Finally, pseudo real-time approaches have been presented where a phased array is continuously rotated internally around its central image axis (fig. 1.4c) [Belohlavek et al. 2001; Canals et al. 1999; Djoa et al. 2000]. For both approaches by Belohlavek et al. [2001] and Canals et al. [1999] the rotation direction is periodically alternated to prevent the cables from getting damaged. The acquisition durations are limited. This provides enough data for 3D LV volume quantification, but restricts the volume reconstructions to low frame rates.

continuous rotating

The design by Djoa et al. [2000] has been extended to harmonic imaging by Voormolen et al. [2006]. A prototype of this transducer has been used in all studies in this thesis. It features a phased array that is continuously rotated at high speed in one direction, employing a slipring construction (fig. 1.5). This fast rotating ultrasound (FRU) transducer (see section 4.1.1) allows long acquisitions up to 10 seconds, which are used for pseudo real-time 3D volumetric reconstructions

(fig. 1.2c). Temporal resolutions up to 25 phases per cardiac cycle are achieved, independent of the patient's heart rate.



Figure 1.5: The fast rotating ultrasound (FRU) transducer

1.2.2.3 | Real-time 3D imaging

Truly real-time 3D echocardiography has been realized originally at Duke University. Von Ramm et al. were the first to build a matrix transducer for real-time 3D imaging [Smith et al. 1991; von Ramm et al. 1991]. Subsequent developments by this group led to the first commercially available 3D phased-array system at the end of the 90's (Volumetrics Model 1, Volumetrics Medical Imaging, Durham, NC).

*current 3DE
systems*

Second generation matrix transducers were introduced by Philips Medical Systems (Best, the Netherlands) and later by General Electric (Milwaukee, Wisconsin, USA). The Philips Sonos 7500 scanner with a X4 xMatrix transducer is capable of live imaging a narrow volume of $25^\circ \times 90^\circ$ at a frame rate of 25 Hz. Full volume imaging is achieved by stitching four narrow image sectors, acquired from seven consecutive beats, together into one volume. It has been succeeded by their Sonos iE33 system with its X3-1 transducer, which shortens full volume imaging to only four cardiac cycles. The same approach is followed by General Electric with their Vivid 7 scanner and its 3V transducer. Recently, also Toshiba and Siemens announced their 3DE systems, of which the latter claims to be able to do real-time imaging of a $90^\circ \times 90^\circ$ volume at 20 Hz, eliminating the need of any ECG gating in the 3D acquisition.

*matrix
advantages*

A major advantage of the real-time scanners is their ability to show live 3D renderings, while acquiring data. Pseudo real-time solutions using mechanically rotated phased-array transducers (such as the FRU transducer) rely on off-line anal-

ysis of the acquired data to achieve a 3D rendering.

Advantages of the FRU transducer over matrix transducers are its better image quality in the 2D image frames and its cost-effectiveness. Furthermore, the FRU transducer allows reliable quantitative analysis based on single-beat data, featuring 6 to 8 2D images per cardiac phase (if using 16-20 phases per cycle) [Voormolen et al. 2007].

*FRU
advantages*

Clinical application | 1.2.3

3D echocardiography has some clear advantages over 2DE in the clinical environment.

At first, the standard 2DE apical views (2-chamber, 4-chamber, long-axis view) can be acquired at once, reducing the acquisition time. Also, 3DE does not suffer from foreshortening because anatomical plane selection can be done off-line, resulting in true standard views. Furthermore, any plane can be visualized off-line.

Secondly, the full geometry of the left ventricle can be imaged. This eliminates the need of making assumptions about the LV geometry in quantitative analyses. This allows more accurate estimation of clinically important parameters such as full cycle LV volume and its derived parameters (ejection fraction, stroke volume, cardiac output, etc.) [Jenkins et al. 2004]. Also, better insight in the LV and the valve geometry is given through 3D renderings of the left ventricle, including better visualization of the wall motion and possible abnormalities. This makes 3DE also a very promising successor for routine stress echocardiography, since regional wall motion abnormalities can be located much more accurately.

Stress echocardiography | 1.2.3.1

Another application of 3D echocardiography involves stress echocardiography. In stress echocardiography patients are examined at different stages of physical or pharmacologically induced stress to visualize regional wall motion abnormalities as a result of myocardial ischemia. 2D stress echocardiography (2DSE) has become a well established tool for identification of patients with coronary artery disease [Armstrong and Zoghbi 2005; Geleijnse et al. 1997]. 3D stress echocardiography (3DSE) has shown to improve on several limitations of 2DSE, such as better anatomical plane selection for comparison of identical wall segments in the different stress stages. Current limitations of 3DSE however, include serious drop outs in the LV lateral wall from rib shadowing, limited temporal resolution and stitching artifacts as a results of volume stitching. All these limitations are expected to be handled by technical developments in 3DE, resulting in smaller transducer footprints,

*3D stress echo-
cardiography*

larger bandwidth transducers suitable for harmonic imaging and higher temporal resolution as a result of parallel beam forming.

1.2.3.2 | Limitations

limitations Current limitations of 3DE are its slightly compromised image quality if compared to 2DE. 3DE spatial image resolution is lower and sensitivity of matrix transducers is still inferior to 2DE. This results in typical image artifacts, such as serious drop outs in the lateral wall region. The temporal resolution of 3DE is also much lower than that of 2DE, which limits the use to patients with relatively low heart rate if reliable estimation of the volume-time curve is needed. Furthermore, high costs are associated with 3DE, which makes 3DE much less commonly available. Also, 3D imaging of the heart has shown to require adequate training of the sonographer.

Most of the limitations mentioned above are expected to be tackled soon, as we currently see rapid developments in transducer design that allow real-time imaging of larger volumes and higher frame rates, with better image quality and resolution. It is to be expected that eventually 3DE will replace 2DE in clinical routine examinations.

1.2.4 | 3DE vs. other modalities

Because of its cost-effectiveness, echocardiography is an attractive imaging modality. It is usually widely available and ultrasound devices can be made portable, allowing bed-side imaging. Furthermore, ultrasound imaging is non-invasive and does not employ ionizing radiation. No adverse biological effect has been reported so far, provided that guidelines for use of diagnostic ultrasound are respected [Barnett et al. 2000].

In diagnosing cardiovascular disease, various other imaging modalities are available, each with its specific strengths and limitations. We will discuss the most common other modalities in image guided diagnosis of cardiovascular disease.

1.2.4.1 | Cardiac magnetic resonance imaging

golden standard Magnetic resonance imaging (MRI) is an important imaging technique that allows non-invasive 3D imaging of the human body. MRI is especially suitable for imaging of soft tissues. The in-plane spatial resolution of current 1.5T scanners is approximately 1.5×1.5 mm, which is comparable to 3D echocardiography. The high temporal resolution of more than 30 frames per second, when using steady-state free precession sequences, combined with the high contrast resolution, makes MRI very suitable for imaging of the heart. Therefore, MRI is accepted as the golden




Figure 1.6: Example images from different patients showing several modalities for imaging the left ventricle. *a)* A short axis slice from a 3D ultrasound image *b)* A short axis MR image *c)* A short axis CT image.

standard for assessment of left ventricular function (fig. 1.6b).

Just as for ultrasound, various different pulse sequences can be used, to target the imaging protocol to specific tissue types or physiological processes. This makes MRI suitable for one-stop-shop approaches, in which various clinically important parameters can be assessed in a single, although possibly time-consuming, scan session. Such a one-stop-shop session might include assessment of LV and RV (regional) functional parameters, myocardial perfusion imaging and late enhancement MRI for localization of infarcted regions and assessment of myocardial viability.

one-stop-shop

Despite its favorable image quality and its harmless nature, MRI is not used as the main imaging modality in cardiology because of the time needed for imaging, the high costs associated with the systems and therefore, their limited availability. Furthermore, the high magnetic fields and powerful radio frequency pulses prevent its use on patients who have metal implants and cardiac pacemakers.

limitations

If compared to 3D echocardiography, several issues have to be kept in mind. The high resolution in the 2D MRI (typically short-axis) images are compromised with a poor through-plane resolution of up to 10 mm. In assessment of left ventricle function this is especially an issue when it comes to defining the base of the left ventricle. This, together with the partial volume effect [Lorenz et al. 1999], different appearance of papillary muscles and trabeculae and different (semi-automated) analysis methods hamper the direct comparison of assessments by MRI and 3DE [Voormolen and Danilouchkine 2007].

MRI vs. 3DE

1.2.4.2 | Computed tomography

Computed tomography (CT) is an imaging technique based on X-ray imaging. A large number of 2D images is taken around a fixed rotation axis to reconstruct a 3D volumetric image. Current multi-slice CT (MSCT) scanners employ up to hundreds of detector rings to acquire the 3D volume faster and allow a higher temporal and spatial resolution. The high temporal resolution opens doors for functional assessment of the heart (fig. 1.6c), and for coronary angiography and perfusion studies.

The main drawback of MSCT, in comparison to MRI and 3DE, is the radiation exposure that is associated with the acquisitions. This will limit the use of MSCT for (global) functional assessment of the left ventricle. Coronary angiography and calcium scoring might be of more interest, although recent studies show that if it comes to ruling out coronary artery disease (CAD), the high negative predictive value of MSCT is compromised with a moderate positive predictive value [Schuijff and Bax 2008].

1.2.4.3 | Nuclear imaging

SPECT Single photon emission computed tomography (SPECT) is a 3D technique that images the distribution of an, intravenously injected, radiopharmaceutical in the body. It can be used to assess myocardial perfusion during different stages of physical or pharmacologically induced stress [Bax et al. 2000; Corbett and Ficaro 1999]. The resolution of this technique is limited and morphological information about the left ventricle is poor. But differences in myocardial perfusion between the stress and rest stage reveal valuable diagnostic information about infarcted regions [Corbett and Ficaro 1999]. Integration with CT allows registration of the perfusion data to the morphological CT images.

PET Positron emission tomography (PET), is a similar technique that also images the distribution of a radiopharmaceutical that indicates tissue metabolic activity. It can be used to detect coronary artery disease with high sensitivity and specificity [Williams 1994]. PET scanners can, like SPECT, be integrated with MR or CT. However, PET has a limited role in routine diagnose of myocardial defects, because of the high costs associated with the production of the necessary radionuclides.

1.3 | Digital image analysis

Images are everywhere. Recently, digital cameras have become so widely available that they did not only replace traditional analog cameras but also got integrated in

mobile phones, PDAs, laptops and other handheld devices. In medicine, a similar progression has been going on. Traditional analog X-ray systems are being replaced by their digital counterparts and echocardiography tapes are being replaced by CD's, DVD's and hard disks. Other complicated digital imaging techniques have become available thanks to developments in computer science (MRI, CT). Nowadays lots of different imaging modalities are available to assist in diagnosis. So many images are acquired, that automation in acquiring, reconstructing, enhancing and analyzing them has become essential.

General purpose of automated image analysis | 1.3.1

Of course, images are acquired primarily for visual inspection and to get insight into the anatomy and physiology of the organ of interest. But there is a growing demand towards techniques that can automatically analyze all these images and derive as much quantitative information as possible from them, such that image guided diagnosis and treatment is brought to a higher level and made more efficient and reproducible.

Automated analysis techniques in general aim at decreasing interobserver variability by ruling out random variability and judgment differences of human experts, making the analysis more reproducible and comparable among institutions. This may lead to a high degree of standardization, which eases the design of protocols and decision making. A high reproducibility is important if the progression of a certain disease is monitored over time. Also, the speed of the analysis can be improved and thereby the labour intensiveness and the costs of the diagnosis or treatment is reduced.

*decreasing
interobserver
variability*

reduce costs

Image analysis improves diagnosis and treatment by quantification of observations, either with or without human intervention in the analysis. Quantification gives more insight in the decision process and can ultimately lead to automated diagnosis, assisting the physician.

quantification

Since manual analysis of 3D echocardiography is cumbersome and very labour intensive, we aim at automating the analysis of 3D echocardiography. We try to automatically quantify the functioning of the left ventricle, to reduce interobserver variability and improve the reproducibility of the quantification results.

This work encompasses the reconstruction of the 3D image over time, from a sequence of 2D images for proper visual inspection of the 3D (plus time) data. Also, this image reconstruction acts as a preprocessing step to allow generalized algorithms for analysis of 3D (plus time) images. For quantitative functional analysis we aim at automated tracking of feature points and structures, to visualize and quantify change in position, size or shape, and orientation of these elements over

*3D image
reconstruction*

segmentation time. This functional analysis also encompasses image segmentation, the automated detection of structures in the 3D image sequences. In these tracking and segmentation algorithms knowledge about image acquisition, the specific patient or the patient population and the targeted structure (or organ) is used to optimize its performance. In the next section we will discuss several important issues in automating these procedures in echocardiography.

1.3.2 | Automated analysis of echocardiography

Unlike other tomographic modalities such as CT and MRI, ultrasound images are hard to interpret, since there is no simple physical relation between the observed image intensity and the imaged medium. Interpreting 3DE is therefore not only a challenge for the untrained human eye, but even more for automated image processing techniques. In the automated analysis of 3DE we have to deal with several ultrasound specific image characteristics.

1.3.2.1 | Image characteristics

speckle Ultrasound image gray values are a result of a summation of sound reflections and scattering, resulting in a combination of interference patterns, called speckle patterns. These patterns give a granular appearance to the image. Differences in imaged media or tissues are observed through differences in these speckle patterns and their intensities. Therefore, transitions between different types of tissue need not render a clear edge in the image, but might show only subtle differences. This granular appearance of the image might challenge the interpretation of the image, but can be of great value when imaging translations and deformations and make ultrasound very suitable for tracking approaches.

position dependency The object appearance is also position dependent in ultrasound imaging. The signal depends on the depth and the objects in the line of sight. Acoustically dense structures might drop a shadow on regions further away from the transducer. Attenuation can be compensated manually using time gain compensation (TGC) while acquiring, but lower signal-to-noise ratios in distant image regions can, of course, not be compensated. The angle-of-incidence of the ultrasound beam influences the reflection and scattering and because of the fanlike acquisition of subsequent beams, echocardiographic images are also highly anisotropic.

Several other image artifacts can be caused by side and grating lobes, reverberations, aberration and noise. Some of these artifacts might be reduced by using harmonic imaging [Duck 2002].

Considerations for automated analysis | 1.3.2.2

For as long as echocardiographic images have been made, also attempts to automatic analysis strategies have been reported. An overview of quantitative methods in 2D echocardiography has been given in Bosch [2007]. Noble and Boukerroui [2006] published a general review of ultrasound image segmentation, also for non-cardiac applications. We will shortly discuss the main considerations when developing an automated analysis approach, relevant to the subject of this thesis.

Ultrasound image appearance is characterized by its granular appearance from speckles and its artifacts as described in the previous section. A great advantage of echocardiographic imaging, at least in 2D, is its high frame rate. These three aspects (speckle, temporal information, typical artifacts) should be considered when designing an analysis technique, whenever possible.

Image speckle can be used as a local image feature. On a small scale speckle serves as a distinct image feature that can be exploited by tracking approaches, as long as object movement is small, relative to the speckle size. This is often employed in 2D echocardiography [Behar et al. 2004; DeCara et al. 2005]. It should be noted however, that speckle patterns depend on the imaging system and that they can change considerably as a result of deformation of the tissue or change in orientation with respect to the transducer. Despite these limitations, texture characterization has been successful in various ultrasound applications [Christodoulou et al. 2003; Sivaramakrishna et al. 2002; Yoshida et al. 2003]. On a larger scale speckle might be an undesirable feature, resulting in a non-Gaussian gray value distribution. Various models have been presented that model the gray value distribution in ultrasound images, which can be incorporated into the detection method [Mignotte and Meunier 2001]. Alternatively, a preprocessing step is often applied, which removes speckle from the image and possibly also aims at transforming the gray value distribution into a Gaussian distribution, such that more general image processing approaches that rely on this property can be applied [Tauber et al. 2004; Xiao et al. 2004; Yu and Acton 2002].

speckle as a feature

speckle suppression

An important source of information comes from the temporal domain. In this domain we can identify static image artifacts, for example as a result of rib shadowing or near-field clutter, and remove noise, such as in the far field. Apart from identifying image artifacts, the temporal information provides most of the functional information we want to extract from an image sequence. The temporal domain can be exploited as multiple observations of a static scene. In this way the object's dynamics are just observed without enforcing any constraints on the dynamics. More robust detection solutions model the object's dynamics to constrain the motion and deformation to expected behavior, as has been elegantly employed in Friedland and Adam [1989] and Comaniciu et al. [2004].

time domain

Because of the typical artifacts that are present in ultrasound images, methods that are solely based on local image features are prone to fail. Image artifacts should be actively detected based on regional spatiotemporal image information, such that a reliability measure can be integrated into the detection [Zhou et al. 2005]. Alternatively, higher level knowledge about the object to be tracked or segmented can be incorporated into the method. This information can be provided by the user, but is ultimately integrated into a model. Various knowledge or model based techniques can be applied to deal with these typical artifacts. Knowledge can be integrated by using some simplified mathematical model to represent the shape of the object, for example based on geometrical assumptions. But also the expected image intensities can be modeled as such, as well as temporal behavior of the object. Various methods integrate prior information about shape and texture [Montagnat and Delingette 2000; Xie et al. 2005]. All of these object properties can also be learned from a training population, for example using a neural network approach as in Binder et al. [1999]. Another concept that is capable of modeling such properties are statistical models, for example active appearance models. These models have been successfully applied to detection of endocardial contours in 2D echocardiographic sequences [Bosch et al. 2002].

image artifacts

*model based
detection*

In chapter 6 we will explore the application of active appearance models for automated segmentation of the left ventricle in 3D echocardiography. In this chapter we also review most of the approaches that have been presented for automated segmentation in 3D echocardiography.

1.4 | Outline of this thesis

Automated estimation of left ventricular volume has been the subject of research for many years. The recent developments in real-time 3D echocardiography have made the assessment of full cycle left ventricular 3D images feasible as a quick, non-invasive, relatively cheap and therefore potentially widely available technique. Manual analysis of the 3D time series of these data sets, however, is cumbersome and subjective, and therefore causes relatively high inter- and intraobserver variabilities in quantifying global left ventricular function. This limits the application in large, inter-institution clinical trials and hampers the value in diagnosis.

We have been challenged by the possibility of assessing global left ventricular function by real-time 3D echocardiography and by the success of previous model-based automated detection attempts to estimate the desired parameters from 2D echocardiography. This has been an inspiration for further improvement of auto-

mented assessment of these important clinical parameters using the opportunities that are provided by the recent innovations in ultrasound imaging.

The main goal of our work therefore is the automation of left ventricular volume quantification using model-based segmentation in 3D echocardiography. This will improve the ease of use of real-time 3D echocardiography for assessment of important clinical parameters for diagnosis of left ventricular function. It will save costly time in analysis of the increasing number of clinical assessments and improve the availability of such parameters in daily clinical practice, with high accuracy and precision, thereby allowing better diagnoses. In this work we have investigated the use of the fast rotating ultrasound (FRU) transducer for real-time 3D echocardiography, which combines advantages of conventional 2D echocardiography with the hugely improved insight given by 3D echocardiography, keeping the general applicability of the developed image analysis techniques in mind. *the main goal*

As a first step to a supervised, fully automatic technique we have worked on a semi-automatic solution for left ventricular volume quantification, which detects full-cycle volumes using limited user interaction based on 2D endocardial border detection in a four-dimensional framework. The interactive nature of this technique allows rapid segmentation of the left ventricle with high accuracy. This is a requirement for the development of a supervised fully automatic technique. The challenges in the extension of previous work to application in higher dimensions and the evaluation of this method have been described in chapter 2.

Chapter 3 studies an important element in full cycle left ventricular volume measurement: tracking the position of the mitral annulus. A substantial time reduction in these full cycle analyses could be achieved by automatic tracking of this quickly displacing anatomical structure. We have studied this problem in 2D echocardiography, with possible application in 3D echocardiography. We present a tracker for 2D structures over time, assuring a time-continuous solution for the mitral annular movement.

For the endocardial detection using native 3D or 4D imaging techniques we have studied the interpolation of the sequence of 2D images acquired in 3D within several consecutive cardiac cycles using the FRU transducer. This work, which deals with multi-beat fusion and the sparse, irregular distribution of the data, is described in chapter 4. An improved method is presented for interpolation of these numerous 2D images from consecutive cardiac cycles into one high resolution 4D cycle.

These high resolution reconstructions of the left ventricle allow native 3D or 4D model-based segmentation approaches for the detection of left ventricular volume. A common problem in the use of model-based segmentation techniques is the initialization of such models in a new data set. We have investigated the rarely studied subject of detection of the main orientation of the left ventricle in 3D acquisitions

for initialization purposes. Chapter 5 discusses a technique for automated detection of the left ventricular long axis and the mitral valve plane. Knowledge about the position of these structures may improve model-based segmentation techniques significantly, since the performance of these techniques often depends on the quality of its initialization.

Chapter 6 presents a fully automatic segmentation technique for the estimation of left ventricular volume based on active appearance models. In this chapter we discuss the adaptation of these models to 3D echocardiography and explore the applicability of active appearance models with different matching approaches.

Finally, we conclude this thesis in chapter 7 and discuss the presented work with recommendations for future research in this direction.

Semi-automatic endocardial border detection for left ventricular volume estimation in 3D echocardiography

2

WE PROPOSE A SEMI-AUTOMATIC endocardial border detection method for LV volume estimation in 3D time series of cardiac ultrasound data. It is based on pattern matching and dynamic programming techniques and operates on 2D slices of the 4D data requiring minimal user-interaction.

We evaluated on data acquired with the fast rotating ultrasound (FRU) transducer: a linear phased array transducer rotated at high speed around its image axis, generating high quality 2D images of the heart. We automatically select a subset of 2D images at typically 10 rotation angles and 16 cardiac phases. From four manually drawn contours a 4D shape model and a 4D edge pattern model is derived. Pattern matching and dynamic programming is applied to detect the contours automatically. The method allows easy corrections in the detected 2D contours, to iteratively achieve more accurate models and improved detections.

An evaluation of this method on FRU data against MRI was done for full cycle LV volumes on 10 patients. Good correlations were found against MRI volumes ($r = 0.94$, $y = 0.72x + 30.3$, a difference of 9.6 ± 17.4 ml (mean $\pm$ standard deviation)) and a low interobserver variability for 3DE ($r = 0.94$, $y = 1.11x - 16.8$, difference of 1.4 ± 14.2 ml). On average only 2.8 corrections per patient were needed (in a total of 160 images). Although the method shows good correlations with MRI without corrections, applying these corrections can make significant improvements.

This chapter has been derived from (© 2005 SPIE):

Semi-automatic border detection method for left ventricular volume estimation in 4D ultrasound data. M. van Stralen, J.G. Bosch, M.M. Voormolen, G. van Burken, B.J. Krenning, R.J.M. van Geuns, E. Angelié, R.J. van der Geest, C.T. Lancee, N. de Jong, J.H.C. Reiber. Proc SPIE Med Imaging 2005; 5747: 1457-1467.

2.1 | Introduction

3D echocardiography

For diagnosis of cardiovascular diseases, the volume and ejection fraction of the left heart chamber are important clinical parameters. 3D echocardiography (3DE) offers good opportunities to visualize the whole left ventricle (LV) over the complete cardiac cycle. 3D echocardiography is non-invasive, relatively cheap, flexible in use and capable of accurate volume measurements [Jenkins et al. 2004; Nosir et al. 1999]. New, fast 3D ultrasound imaging devices are entering the market and have the potential of allowing such measurements rapidly, reliably and in a user-friendly way - provided that a suitable automated analysis is available. Manual segmentation of the large data sets is very cumbersome and suffers from inconsistencies and high variability. On the other hand, the human expert's interpretation and intervention in the detection is often essential for good results. Therefore a semi-automatic segmentation approach seems most suitable.

2.1.1 | Other approaches

Some methods for segmentation of 3D echocardiographic images have been published. Angelini et al. [2001] have reported on a wavelet-based approach for 4D echocardiographic image enhancement followed by a segmentation of the left ventricle using snakes. Corsi et al. [2002] presented a level-set based semi-automatic method. Montagnat and Delingette [2000] used a 2-simplex mesh and a feature detection based on a simple cylindrical gradient filter. Sanchez-Ortiz et al. [2002] used multi-scale fuzzy clustering for a rough segmentation in 2D longitudinal slices. B-splines are used for 3D surface fitting in each time frame. These methods have not been validated successfully on a reasonable data set. The most practical approach is described by Kühl et al. [2004]. It uses active surfaces that are controlled by difference-of-boxes operators applied to averages and variances of the luminance. This technique is implemented in a commercially available workstation (4D LV Analysis, TomTec, Unterschleißheim, Germany). The general experience is that this technique requires much initialization and corrections, and a consistent segmentation is still hard to reach. Another commercial development has been presented recently: QLAB (Philips Medical Systems, Best, the Netherlands). This package provides on- and offline 3D quantification tools. However, technical details or clinical evaluations of these methods have not been reported yet.

our approach

We present a semi-automatic endocardial border detection method for left ventricular volume estimation in time series of 3D cardiac ultrasound data. Our method is based on pattern matching and dynamic programming techniques and com-



Figure 2.1: The fast rotating ultrasound (FRU) transducer

bins continuity, robustness and accuracy in 2D cross sections with the spatial and temporal continuity of the 3D plus time (3D+T) data. It aims at optimally using a limited amount of user interaction (capturing essential information on shape and edge patterns according to the user's interpretation of the ultrasound data) to attain a fast, consistent and precise segmentation of the left ventricle.

Despite the fact that this method is optimized for data of the fast rotating ultrasound transducer (see below), the algorithm can be easily adapted to data of other image acquisition systems, for example 3D+T voxel sets. The detection will then be performed in 2D slices through the LV long axis.

*general
applicability*

Fast rotating ultrasound transducer | 2.1.2

We performed this study on a special type of image data acquired with a new device: the fast rotating ultrasound (FRU) transducer (fig. 2.1). The transducer has been developed by the Department of Experimental Echocardiography of the Erasmus MC, the Netherlands [Djoa et al. 2000; Voormolen et al. 2002]. It contains a linear phased array transducer that is continuously rotated around its image axis at very high speed, up to 480 rotations per minute (rpm), while acquiring 2D images. A typical data set is generated during 10 seconds at 360 rpm and 100 frames per second (fps). The images of the left ventricle are acquired with the transducer placed in apical position, with the transducer's rotation axis more or less aligned with the LV long axis. The analysis assumes that the rotation axis lies within the LV lumen and inside the mitral ring.

*linear phased
array
transducer*

An important advantage of this transducer is that it can be used with any ultra-




Figure 2.2: A sequence of seven consecutive FRU images with curved image planes

*usage with any
ultrasound
machine* sound machine, since a conventional phased array transducer is used. It also acquires relatively high quality 2D images, compared to matrix array transducers used for real-time 3D echocardiography. Furthermore, no ECG triggering is applied, just an ECG-registration for offline analysis, which allows quick acquisitions.

*acquisition is
not ECG
triggered* As a consequence of the very high continuous rotation speed, the images have a curved image plane (fig. 2.2). During the acquisition, the probe rotates about 22° per image with the typical settings given above. The combination of these curved image planes, and the fact that the acquisition is not triggered by or synchronized to the ECG signal, results in an irregular distribution over the 3D plus time (3D+T) space. A single cardiac cycle in general is not sufficient for adequate coverage of the whole 3D+T space; therefore, multiple consecutive heart cycles are merged. The cardiac phase for each image is computed offline using detected R-peaks in the ECG [Engelse and Zeelenberg 1979]. From the total set of ± 1000 2D images, a subset of images with a regular coverage of the 3D+T space is selected automatically. We perform analysis on the images in this subset. The data is also suitable for the generation of a time series of 3D voxel sets.

Methods | 2.2

Frame selection | 2.2.1

To achieve adequate coverage of the whole 3D+T space, multiple consecutive cardiac cycles are merged and an optimal subset S of the total set of frames T is selected (fig. 2.3). This subset is an optimal fit of the frames on a chosen $A \times P$ matrix of A equidistant rotation angles and P cardiac phases, minimizing the total deviation in rotation angle and cardiac phase. Moreover, the variation in acquisition time over the subset is minimized to limit possible motion artifacts. The constraints are translated into the following cost functions that will be minimized over the total subset S ,

limit motion artifacts

$$S = \bigcup_{i=1}^A \bigcup_{j=1}^P (\operatorname{argmin}_{b \in C_{i,j}} (c_{\text{angle}}(\alpha_b, i) + c_{\text{phase}}(p_b, j) + c_{\text{time}}(t_b))) \quad (2.1)$$

$$c_{\text{angle}}(\alpha, i) = k_1 |\alpha_{\text{target}}(i) - \alpha|$$

$$c_{\text{phase}}(p, j) = k_2 |p_{\text{target}}(j) - p|$$

$$c_{\text{time}}(t) = k_3 |t_S - t|$$

$C_{i,j}$ is the set of candidate images for angle $\#i$ and phase $\#j$. c_{angle} and c_{phase} describe the costs of selecting an image b with angle α_b and phase p_b for a chosen α_{target} and p_{target} . k_1 , k_2 and k_3 are weighting coefficients (typically equal). Since the cost c_{time} is dependent on t_S (the average acquisition time of the subset itself), the minimization of the costs of set S is achieved in an iterative manner.

iterative frame selection optimization

Border detection approach | 2.2.2

We base our method on the knowledge that the edge patterns of the endocardial border can be complex, very different from patient to patient and even between regions within an image set. The border position need not correspond to a strong edge and may be only definable from 'circumstantial evidence' as identified by an expert observer. Rather than applying artificial, idealized edge models or templates derived from a large training set, we propose a tracking approach based on edge templates extracted from the user-defined initial borders in the patient's own images.

patient specific edge template

The method is based on the following continuity assumptions (in order of strength):

- (a) border continuity within separate 2D slices of the left ventricle;




Figure 2.3: Selected subset of 2D FRU images in 16 cardiac phases and 10 rotation angles. Contours are manually drawn in the highlighted images.

(b) spatial continuity of shape and gray value edge patterns over the LV surface in 3D;

(c) temporal and cyclic motion continuity of the endocardium.

For the FRU transducer, within the original 2D images, both spatial and temporal distances between neighboring samples are smaller than towards adjacent images in angle and phase; therefore, border continuity is supposed to be strongest here.

manual initialization The method is initialized from four manually drawn contours, taken from two roughly perpendicular views (more or less corresponding to two- and four-chamber cross sections) in two phases: end diastole (ED) and end systole (ES). These are used to initialize a model for the edge patterns near the 3D LV surface over time and a 3D shape model of the LV endocardial surface over the entire cardiac cycle. Both models are inherently 4-dimensional and can be polled at any spatial position and cardiac phase.

method overview The actual border detection takes place in individual 2D images from the selected subset and is an extension of an approach for 2D+T sequences earlier developed by Bosch et al. [1998]. For each image $b \in S$ (of cardiac phase p_b and rotation angle α_b), an estimation of the border shape is derived by intersecting the 3D shape model at phase p_b by the (curved) image plane for angle α_b . The edge templates are also interpolated for the desired p_b and α_b . In the 2D image, a neighborhood of the estimated shape is resampled along lines perpendicular to the shape estimate. Using a template matching with the local edge templates, the similarity of

each candidate edge point to the template is calculated. Dynamic programming is applied to find an optimal continuous border within the restrictions posed by the 3D model. In this way, the 3D+T surface and edge pattern models guard the (looser) spatial and temporal consistency of the detection, while the dynamic programming approach supplies a continuous and optimal detection locally. The set of detected contours describes the 3D endocardial surface over the whole cardiac cycle from which LV volumes, ejection fraction and other valuable parameters can be computed.

3D surface models | 2.2.3

Definition | 2.2.3.1

As said, for two cardiac phases (ED and ES) a 3D surface model of the LV endocardium is constructed from two almost perpendicular contours. During the acquisition the rotation axis is more or less aligned with the long axis (LAX) of the left ventricle, but in practice there may be a considerable mismatch (fig. 2.4b). This implies that the two image planes do not contain the true apex of the left ventricle, and estimating the position and shape of the true apex (and the LV long axis) is a non-trivial issue. The local long axes in the 2D manually drawn contours are defined as the lines through the midpoint of the mitral valve (MV) and center of gravity of the upper 10% of the contour area. We estimate the 3D LV long axis from the local long axes by computing the intersection of the planes perpendicular to these images through the local long axis in the image.

*manual
contours miss
apex*

The endocardial surface is estimated by expressing the two contours in a cylindrical coordinate system with respect to the estimated LV long axis. Intersection points of these contours are found with a stack of planes perpendicular to the long axis (short-axis planes). Within each short-axis plane, a closed contour is found by interpolating between the intersection points; for this, the radial coordinate component r is interpolated over the angle between the intersection points (see section 2.2.3.2 for details). This gives a natural approximation of the ellipsoidal shape of the left ventricle. Since the two image planes generally do not intersect the real apex, the apical cap of the LV surface cannot be estimated simply from the two manually drawn contours, as shown in fig. 2.4b. Therefore, near the 3D apex we use a spherical coordinate system oriented around the LV long axis, centered at $3/4^{\text{th}}$ of its length. The surface is estimated by interpolating the radial component over the elevation angle for multiple rotation angles, using the interpolation method described in the next section. A contour estimate for any 2D image at a given rotation angle and cardiac phase can be made by intersecting its curved image plane with

*shape models
in ED and ES*




Figure 2.4: *a*) The interpolation of the endocardial surface in a cylindrical coordinate system oriented around the LV long axis (LAX). *b*) 3D surface model. The LAX estimate (dotted) and the rotation axis (dashed) are shown, together with the reconstruction of the apex by spline interpolation (light gray) from two manually drawn contours (solid black). *c*) The extraction of a stylized edge pattern from an image with a manually drawn contour.

the 3D contour models in ED and ES and then linearly interpolating between the two resulting '2D' contours over cardiac phase to get the contour estimate at the desired cardiac phase.

2.2.3.2 | Surface interpolation/fitting

Fitting a smooth contour through all available intersection points in a short-axis plane is not always possible. Inconsistencies can occur in the set of input contours used for the interpolation of the endocardial surface interpolation. They can be caused by inconsistent manual tracing or by inconsistent image data. The latter can be caused by substantial differences in cardiac phase between the images or by inter-beat variation. For the generation of a smooth endocardial surface, we developed a fitting algorithm that can handle these inconsistencies.

dynamic programming The algorithm is dynamic programming based. Dynamic programming (DP) [Sonka et al. 1999] is a well known graph search technique that finds the optimal path through a rectangular array of nodes (the path with the lowest sum of costs) out of all possible connective paths in an effective manner by calculating lowest cumulative costs for consecutive layers (lines) while keeping track of the partial optimal paths. Backtracking from the node with lowest cumulative cost in the last

layer delivers the overall optimal path. A connective path contains exactly one node per line and the positions on consecutive lines cannot differ more than a predefined side step size.

Using this technique, the method fits an optimal curve through a set of possibly inconsistent intersection points. It allows the assignment of reliabilities to each point. Also, the curvature can be controlled through parameters in the dynamic programming algorithm and the probability distribution computation, which is explained below.

control the curvature

The curve is found through the set of intersection points $i \in I$ with corresponding reliabilities p_i . The nodes in the DP array of size $A \times R$, represent points in (α, r) -space. Finding the path with the minimum costs solves the fitting problem. The costs of each node are represented by the cost function C ,

$$C(n) = -\log(P^*(n)), \quad (2.2)$$

where $P^*(n)$ is the normalized probability that node n represents a point on the endocardial border. The normalization is performed within each layer of the DP graph, such that the probabilities within each layer sum up to one. The probability $P(n)$ of node n being part of the endocardial border is inversely related to the angular and radial distance to the intersection points, δ_α and δ_r , and is defined as,

$$P(n) = p_0 + \sum_{i \in I} p_i G(\delta_r, \sigma) \quad (2.3)$$

$$\sigma = c_1(\delta_\alpha + c_2)^{c_3} \quad (2.4)$$

The probability distribution is Gaussian (G) in the radial direction (within the DP layers), as defined in eqn. 2.3. The width of the Gaussian increases with the angular distance from the input point δ_α (eqn. 2.4), which makes the distribution more flat with increasing angular distance. c_1 , c_2 and c_3 are parameters that influence the curvature and smoothness of the resulting curve, where $c_1 > 0$, $c_2 > 0$ and $c_3 > 1$. An example cost matrix and the resulting curve are shown in fig. 2.5.

Edge pattern model | 2.2.4

The desired edges are tracked over space and time by applying a pattern matching approach with edge templates. These edge patterns are derived from the manually drawn contours and interpolated over the (phase, angle)-space. The image is re-sampled clockwise along the manually drawn contour, on line segments perpendicular to this contour from the inside out. The gray values on these line segments are smoothed and subsampled to form a stylized edge pattern for this contour (fig.

derived from manual contours




Figure 2.5: *a)* The dynamic programming cost matrix. *b)* The resulting curve fit through the input point

2.4c). A typical edge pattern for a single 2D frame is represented by 32 positions along the contour and 5 samples around each edge position.

The interpolation over cardiac phase is performed linearly between the edge patterns in ED and ES. The interpolation over rotation angle is less straightforward. Since the character of the edge pattern is strongly related to the angle between the endocardial border and the ultrasound beam and the distance from the transducer, the pattern changes considerably over the rotation angle, especially when the angle between the rotation axis and LV long axis is substantial. For images with rotation angles opposite ($\pm 180^\circ$) to those with the manually drawn contours, the image appears nearly mirrored and the mirrored (anti-clockwise) edge pattern is used. For angles in between, the edge patterns are linearly interpolated.

2.2.5 | Contour detection

With an edge pattern and initial contour for each image $b \in S$ (of phase p_b and angle α_b), we can now detect the individual endocardial borders (fig. 2.6). In a neighborhood of the initial contour, the image is resampled into an $N \times M$ rectangular array by sampling N points along M scan lines perpendicular to the shape. From the styl-

ized edge pattern for (p_b, α_b) an edge template for each scan line is extracted. For all nodes in the array, the sum of absolute differences with its respective edge template defines the cost of the node. We now use a dynamic programming approach (section 2.2.3.2, [Sonka et al. 1999]) to find the optimal connective path through the array. Smoothness constraints are enforced by applying additive costs for side stepping during cumulative cost calculation. To limit the influence of lines with relatively poor image information, this additive penalty is calculated per line from the statistics of node costs per line with respect to overall cost statistics, such that relatively unreliable lines get higher penalties for side stepping.

*dynamic
programming*




Figure 2.6: Contour detection. *a)* Resampling of the image around the 2D shape estimate. *b)* Edge pattern matching and dynamic programming to detect the optimal contour. *c)* The detected contour

For each phase p_j , the detected contours of all angles α_i together constitute a 3D mesh that describes the endocardial surface. We observe the volume of the left ventricle over the whole cardiac cycle, by calculating the volumes inside the surface meshes of all selected cardiac phases.

Correct and redetect | 2.2.6

In the initial detection the shape and edge pattern models are estimated from only four manually drawn contours. In some cases, this does not provide enough information for the models to detect the endocardial border well in all the images in the subset. Also, the border may be poorly defined in some of the images, which complicates the detection. Therefore the method allows additional corrections in the detected contours in the 2D images. A corrected contour will be treated as an

*iterative
refinements*

additional manual contour and both the edge pattern and shape models will be updated accordingly, achieving a more specific approximation of the actual shape and appearance. This results in a new set of shape and edge pattern estimates, and all remaining images are redetected. Through these easy, iterative refinements, corrections will cumulatively lead to a superior global solution.

2.3 | Results

comparison with MRI We performed a preliminary validation study for this method on a group of 10 subjects with different diagnoses of cardiovascular disease. Full cycle MRI LV volumes on these patients were determined in parallel with the 3DE study, using semi-automatic segmentation tools (MRI-MASS, Medis medical imaging systems, Leiden, the Netherlands) by an independent observer unaware of the 3DE analyses. For the 3DE study, for all patients, subsets of images were created with $P = 16$ phases and $A = 10$ angles. After establishing equivalent tracing conventions, two observers individually analyzed all subsets. We evaluated the semi-automatic segmentation method using only the initial four manually drawn contours, and after applying corrections iteratively.

analysis time Reading and converting the data and the automated selection of the subset took 7 minutes per patient on average. After the drawing of the four contours, the fully automated detection of the other 156 contours took approximately 60 seconds per patient (Pentium IV 2.6GHz, 2GB RAM). Some examples of manual and detected contours are shown in fig. 2.7. Corresponding endocardial surfaces for all phases are shown in fig. 2.8. From the analyses of both observers, interobserver variabilities of full cycle volumes and ejection fractions (EF) were determined, as well as averages that were correlated to full cycle MRI volumes, both for the situation with and without applying corrections.

without corrections Results of 3DE (average of the two observers) vs. MRI are shown in table 2.1 and fig. 2.9a-d. Fig. 2.9a shows the full cycle volumes without corrections. A good correlation of $r = 0.92$ was found between MRI and 3DE volumes without corrections. Regression was $y = 0.675x + 32.9$. In general, ED volumes were underestimated by 3DE, while ES volumes were slightly overestimated. Overall the MRI volumes were slightly larger (14.1 ± 19.6 ml (mean $\pm$ standard deviation)). Similar differences between 3DE and MRI volumes have been reported in many studies and can be attributed to differences in tracing conventions between MRI and 3DE. EF results showed a reasonable overall difference of $6.2 \pm 8.9\%$, but the regression was similarly affected ($y = 0.36x + 25.8\%$, $r = 0.63$). For the 3DE interobserver variabil-




Figure 2.7: Detection examples: frames at different (phase#, angle#) with contours. *left*) 4 frames with manual contours, resp. 1: ED 2c (1,1), 2: ES 2c (6,1), 3: ED 4c (1,3), 4: ES 4c (6,3). *right*) 4 frames with detected contours, resp. 5: frame (8,2), 6: (14,5), 7: (4,8), 8: (14,9)




Figure 2.8: Reconstruction of the endocardial border over the full cardiac cycle. The top row shows phases 1 to 8, the bottom row phases 9 to 16 (of a total of 16 phases).

ity, results are presented in table 2.1 and fig. 2.9e. The differences were 5.0 ± 13.6 ml with a regression of $y = 1.08x - 16.0$ ($r = 0.94$).

with
corrections

After applying corrections, the full cycle results slightly improved in comparison to MRI volumes: $r = 0.94$, $y = 0.73x + 30.3$, with a difference of 9.6 ± 17.4 ml (table 2.1, fig. 2.9b,d). This is equally reflected in EF results: $r = 0.64$, $y = 0.36x + 25.8$ with differences of $6.0 \pm 8.8\%$. On average 2.8 corrections were applied per patient (on a total of 160 images). Interobserver variability (fig. 2.9f) increased slightly due to the additional corrections, with a correlation of $r = 0.94$ and $y = 1.11x - 16.8$ (difference of 1.4 ± 14.2 ml).

Table 2.1: Full cycle volumes and correlations of 3DE vs. MRI and Observer 1 vs. Observer 2 (N = 160)

		Volume (ml)	Correlation	Regression
Without corrections	MRI	148 ± 48	0.919	0.675x + 32.9
	3DE	134 ± 36		
	Obs.1	136 ± 34	0.941	1.08x - 16.0
	Obs.2	131 ± 39		
Without corrections	MRI	148 ± 48	0.936	0.728x + 30.3
	3DE	138 ± 38		
	Obs.1	139 ± 35	0.943	1.11x - 16.8
	Obs.2	138 ± 41		

The volumes are expressed as mean ± standard deviation.

2.4 | Discussion

The findings in this study suggest that the semi-automated detection method is a useful tool for quick, semi-automatic detection of LV endocardial borders. However, apparent difficulties in interpreting the 3DE images somewhat obscure the conclusions. Although the different observers can easily reach a satisfactory analysis of the 3DE datasets, the volumetric results show considerable interobserver differences. On inspection, these differences were primarily due to different image




Figure 2.9: All volumes in ml. *a)* MRI vs. 3DE (average of two observers) volumes, without corrections. *b)* MRI vs. 3DE (average of two observers) volumes, with corrections. *c)* Bland-Altman analysis for MRI vs. 3DE, without corrections. *d)* Bland-Altman analysis for MRI vs. 3DE, with corrections. *e)* Interobserver variability (for 3DE), without corrections. *f)* Interobserver variability (for 3DE), with corrections

interpretation and corresponding manual contours, not to differences in detected contours; these consistently followed the observer's interpretation. Clearly, criteria for tracing were not established well enough in this case. This was partly due to the somewhat unconventional cross sections for tracing, partly to the attempt to trace 'similar to MRI criteria' which was not completely successful and meant that we had to deviate from standard ultrasound contour drawing conventions. In some cases, image quality was a factor as well.

Absolute differences in volumes may seem high, but this is partially due to the dilated hearts in the set and the consequent high average volumes (mean MRI ED volume = 187 ml).

tracing conventions It would be useful to extend the study with tracings following normal ultrasound conventions. This could lead to lower observer variabilities, although the comparison to MRI volumes could be more complicated. Comparison between 3DE and MRI volumes was hampered by the different tracing conventions; this resulted in considerable systematic differences, which were also clearly dependent on cardiac phase. Still, overall regression coefficients were high, especially considering the large interobserver variations for ultrasound. This suggests that with proper tracing conventions and/or correction formulae, a high correspondence between MRI and 3DE volumes should be realizable.

Currently, EF measurements by the 3DE method suffer considerably from the systematic differences and variabilities described above. EF measurements should also benefit greatly from improved tracing conventions.

mitral valve tracking Looking at the distribution of volume errors over the full cardiac cycle, a shortcoming of this method becomes clear: the lack of a true mitral valve tracking algorithm. Currently, the method simply assumes the movement of the mitral valve to be linear in systole and diastole, which is a substantial simplification. It is well known that the valve plane motion is directly related to LV volume change, and we clearly observe the effect: volume curves which are a bit too 'linear' in diastole and systole. Despite the possibility to adjust this movement using corrections, still high differences can be observed in early systole and diastole in comparison to MRI volume curves. Extending this method with a mitral valve tracking algorithm is expected to further improve the results. Several approaches for mitral valve tracking have already been presented in literature. We elaborate on this topic in chapter 3.

Conclusions and future work | 2.5

We presented a new semi-automatic endocardial border detection method for 3D+T ultrasound data. This method offers fast and reasonably precise automated border detection with minimal user interaction. The method shows good full cycle results against MRI in the initial detections with only four manually drawn contours. After applying corrections, the results do improve for the individual patients, but in the overall comparison against MRI these improvements do not make a significant difference. This can be addressed to the different interpretation of the 3DE data by the observers in comparison to the MRI data. A satisfying detection in 3DE does not always result in equivalent volumes.

The method can still be improved by including a mitral valve tracking algorithm. Furthermore, better tracing conventions for this type of 3DE data would be helpful for consistent analysis.

Automated tracking of the mitral valve annulus motion in apical echocardiographic images using multidimensional dynamic programming

3

WE DEVELOPED A SEMI-AUTOMATIC method for tracking the mitral valve annulus (MVA) in echocardiographic images, in particular, tracking the septal and the lateral mitral valve hinge points.

The algorithm is based on multidimensional dynamic programming combined with apodized block matching. The method was tested on single-beat apical four chamber image sequences of 20 patients with acute myocardial infarction. The automated tracking results were evaluated by comparing them with the average manual tracking results of two experts. The mitral valve hinge point displacements and the total mitral excursions obtained by the automatic technique agreed well with those obtained manually and outperformed two commonly used tracking methods (forward tracking and minimum tracking).

In conclusion, this novel semi-automatic tracking method is clinically valuable and capable of tracking the MVA motion within the limits of interobserver variability. The technique is robust, even in low frame rate, redigitized VCR images of clinical quality.

This chapter has been derived from © 2007, with permission from Elsevier):

Automated tracking of the mitral valve annulus motion in apical echocardiographic images using multidimensional dynamic programming. S.T. Nevo, M. van Stralen, A.M. Vossepoel, J.H.C. Reiber, N. de Jong, A.F.W. van der Steen and J.G. Bosch. *Ultrasound Med Biol* 2007; 33(9): 1389-1399.

3.1 | Introduction

*mitral valve
annulus
motion*

Analysis of the mitral valve annulus (MVA) movement is useful in the evaluation of global and regional left ventricular (LV) function and is important for diagnosis of mitral annular diseases and LV disorders [Eto et al. 2005; Pai et al. 1991; Willenheimer et al. 1999] In addition, the motion of the MVA plane contributes significantly to the left ventricular volume change over the cardiac cycle and correlates well with LV ejection fraction [Eto et al. 2005].

the purpose

Echocardiography is a non-invasive and low cost technique for evaluating LV function, but it is not an easy modality for interpretation. Therefore, extracting quantitative information from such images is of great clinical importance. The purpose of this work is automatic tracking of the motion of the MVA over the full cardiac cycle; specifically, tracking the mitral valve hinge points (MVHPs). This method can be useful on its own, but we mainly intend it to be part of a semi-automatic border detection approach (chapter 2), to improve the LV volume estimation. Given the septal and the lateral MVHPs in the first frame (fig. 3.1), e.g. by manual annotation, we desire to track them over the full cardiac cycle.




Figure 3.1: The white points represent the septal (left) and the lateral (right) MVHPs in the first frame

related work

Common methods for evaluating the mitral annular motion are M-mode echocardiography [Unser et al. 1989; Willenheimer et al. 1999], two-dimensional (2D) echocardiography [Behar et al. 2004; DeCara et al. 2005; Eto et al. 2005], three-dimensional (3D) echocardiography [Flachskampf et al. 2000] and Doppler tissue

echocardiography [Rabben et al. 2000]. Unser et al. [1989] performed one-dimensional (1D) dynamic programming (DP) on M-mode echocardiograms (see section 1.2.1) for border detection. Rabben et al. [2000] implemented 1D DP on Doppler tissue echocardiograms combined with gray-scale M-mode for contour detection. The dominant drawback of M-mode and Doppler tissue echocardiography is that they detect the MVA motion only along a fixed line (angle dependency). DeCara et al. [2005] results for automated tracking of mitral annular displacement based on speckle tracking. Previous studies [Behar et al. 2004; Eto et al. 2005] suggested block matching combined with forward tracking for tissue tracking using 2D echocardiography. The main drawbacks of forward tracking are the sensitivity to small disturbances and tendency to drift away from the optimal path once an error has occurred.

To overcome such limitations, we propose a 2D echocardiography based tracking technique using multidimensional dynamic programming (MDP) combined with apodized block matching.

Materials and methods | 3.2

Blockmatching | 3.2.1

Block matching is a common technique for motion tracking in ultrasound data [Chen et al. 1995; Leung et al. 2006a]. Given a sequence of 2D echocardiographic images, the user selects the point to be tracked in the reference image (first image) and a region-of-interest (ROI) is defined around it (ROI_1). The displacement of a certain point is estimated by searching for $ROI_t(k, l)$ (with the same size as ROI_1) in the region-of-search (ROS) in the target frame (t) with the most similar gray level pattern to that of ROI_1 (see fig. 3.2). For each shift (k, l) the similarity of the patterns between ROI_1 and $ROI_t(k, l)$ is estimated by a cost function (matching criterion). The shifts (k, l) with respect to ROI_1 that correspond to good values of the cost function are likely displacement estimates [Behar et al. 2004; Leung et al. 2006a]. Three candidates for cost functions were considered:

1. Sum of squared differences (SSD) of gray values between ROI_1 in the reference frame and $ROI_t(k, l)$ in the target frame. This cost function is optimal when it is minimal. The general formula for the SSD [Leung et al. 2006a],




Figure 3.2: The block matching components

combined with the apodization kernel is:

$$SSD(k, l, t) = \sqrt{\sum_{j=j_1-N}^{j_1+N} \sum_{i=i_1-M}^{i_1+M} (|I_1(i, j) - I_t(i+k, j+l)| \cdot W(i, j))^2} \quad (3.1)$$

where $I_1(i, j)$ and $I_t(i, j)$ denote the pixel gray level at position (i, j) of the reference frame and the target frame respectively, $2M+1$ and $2N+1$ are the size of the ROI, (k, l) are the offsets of ROI_t within I_t with respect to I_1 and $W(i, j)$ is an apodization kernel that gives high relevance to central pixels.

2. Sum of absolute differences (SAD) of gray values between ROI_1 and $ROI_t(k, l)$. This cost function is optimal when it is minimal. The general formula for the SAD [Leung et al. 2006a], combined with the apodization kernel is:

$$SAD(k, l, t) = \sum_{j=j_1-N}^{j_1+N} \sum_{i=i_1-M}^{i_1+M} (|I_1(i, j) - I_t(i+k, j+l)| \cdot W(i, j)) \quad (3.2)$$

3. Normalized cross correlation (NCC) of ROI_1 and $ROI_t(k, l)$. This cost function is optimal when it is maximal. The general formula for the NCC [Bohs and

Trahey 1991], combined with the apodization kernel is:

$$NCC(k, l, t) = \frac{\sum_{j=j_1-N}^{j_1+N} \sum_{i=i_1-M}^{i_1+M} \mathcal{I}_1(i, j) \mathcal{I}_t(i, j, k, l)}{\sqrt{\sum_{j=j_1-N}^{j_1+N} \sum_{i=i_1-M}^{i_1+M} \mathcal{I}_1(i, j)^2 \sum_{j=j_1-N}^{j_1+N} \sum_{i=i_1-M}^{i_1+M} \mathcal{I}_t(i, j, k, l)^2}} \quad (3.3)$$

$$\mathcal{I}_1(i, j) = (I_1(i, j) - \bar{I}_1) \cdot W(i, j)$$

$$\mathcal{I}_t(i, j, k, l) = (I_t(i + k, j + l) - \bar{I}_t) \cdot W(i, j)$$

where $\bar{I}_1$ and $\bar{I}_t$ indicate the mean values of ROI_1 and ROI_t , respectively. The NCC has the advantage that it is insensitive to local variations in the mean and standard deviation of the image gray level [Golemati et al. 2003]. For consistency with the other cost functions, we considered the cost function $1 - NCC(k, l, t)$, which is optimal when it is minimal.

Concepts of dynamic programming | 3.2.2

Dynamic programming (DP) is an optimal approach for solving variational problems by finding locally optimal solutions consecutively. DP algorithms provide an efficient way of finding an optimal path through a graph of nodes. In an image processing context, the graph is generally a matrix of image-related cost values ordered with root and leaves from left to right. The maximum distance between two nodes in consecutive columns, the step size, is predefined according to the desired connectivity. The task is to find a connective path through this matrix with minimal total cost. This is achieved by sweeping from left to right through the graph and finding for all nodes the lowest cumulative cost to reach this node from any of its predecessors in the previous column (for which the lowest cumulative cost is already known). Each element in the cumulative cost matrix corresponds to the minimal sum of costs to reach this point from the beginning of the matrix. In the last step, the optimal path is extracted by finding the minimal value in the last column and tracking back through minimal values over the cumulative cost matrix. The result is a global optimal path [Amini et al. 1990; Sonka et al. 1999; Üzümcü et al. 2006].

*graph search
technique*

cost matrix

back tracking

An example of DP search for a 1D path is illustrated in fig. 3.3. The matrix on the left is the cost matrix as derived from an image related cost function (e.g. block matching). Each term in the cost matrix represents a node with a given cost. The matrix on the right is the cumulative cost matrix derived from the original cost matrix when passing through it from left to right. The backtracking begins from the

minimal cumulative cost in the rightmost column and propagates to the left. The optimal path, which connects minimal cumulative costs per column, as calculated from the back-propagation, is shown in red.

multi-
dimensional
dynamic
programming

In most cases, DP is applied in the way described above, resulting in an optimal path through a 2D array. In a tracking problem, the horizontal axis represents the consecutive time steps and the vertical axis represents a certain parameter, for example, 1D displacement. In our tracking problem, the purpose is to track two parameters simultaneously: horizontal and vertical displacement. This can be done by applying multidimensional dynamic programming (MDP) [Üzümcü et al. 2006], 2D DP in our case: finding a path through a stack of slices (a 3D data set), where each slice corresponds to the ROS filled with the values of the 2D cost function that was obtained from the block matching. 2D DP is illustrated in fig. 3.4.

A common variation of DP in tracking is forcing a path through a node. Certain points can be imposed to be on the path by setting the costs of all other nodes of a certain column (in the 1D case) or a certain slice (in the 2D case) to sufficiently high values [Sonka et al. 1999; Üzümcü et al. 2006]. In our case, the user specifies the positions of the MVHPs in the first frame. Next, they are tracked over the cardiac cycle by finding the optimal path that starts at this point.

3.2.3 | Tracking: MDP search combined with apodized block matching

manual
initialization

The coordinates of the septal and the lateral MVHPs are determined in the first frame. After the user manually selects the two points representing the positions of the septal and lateral MVHPs, the automated tracking is started. An image patch around the manually chosen points in the first frame is used as the ROI to be matched. Block matching is applied by computing the cost function for each point in the ROS of each frame. In block matching in medical images, the size of the ROI is crucial. It should be large enough to identify a unique region, but small enough not to suffer too much from deformation of the block content over time. We want




Figure 3.3: Example of 1D DP. *left*) The cost matrix. *right*) The corresponding cumulative cost matrix with the optimal path marked in red




Figure 3.4: Example of 2D DP. Each slice represents the cost matrix at a certain frame. The optimal path is shown in red

to track the motion of the center pixel accurately. However, when using large rectangular blocks, the relative contribution of pixels far away from the center pixel becomes dominant. To limit this effect, we use an apodized block matching (as given by eqn. 3.1 to eqn. 3.3), meaning that within the summation over the ROI, the cost elements are multiplied by a 2D Gaussian kernel that gives higher weight to central elements rather than peripheral ones. The total ROI size is coupled to the width of the 2D Gaussian apodization kernel: it is chosen such that at least 90% of the summed apodization weights are included (the width of the ROI was chosen as 2×2 standard deviations of the Gaussian kernel). A path search using MDP is implemented to find the optimal locations of the septal and lateral MVHPs (nodes in the DP graph) separately over the cardiac cycle. The size of the ROS determines the dimensions of the DP graph. In this way, the horizontal and the vertical displacements of the MVHPs are tracked.

apodization

Parameter selection | 3.2.4

The performance of the method depends on the following parameters:

Width of smoothing filter In the preprocessing phase, the images are smoothed with a Gaussian filter in the spatial and the temporal domain. The variance (the width) of the filter determines the amount of smoothing. A Cartesian approach was followed rather than a polar one since the ROI is always in the same position and therefore Cartesian coordinates correspond well with polar coordinates. Moreover, the degradation of the image quality mainly occurs due to interlacing artifacts, which arise in Cartesian coordinates.

Apodized region-of-interest (ROI) size The ROI size (the size of the template to be matched) is chosen to be four times the standard deviation of the apodization

kernel in each direction. The ROI size determines the amount of information involved in the calculation of the costs.

Region-of-search (ROS) size The size of the search space should be large enough to accommodate the maximal displacement over the cardiac cycle without leading to an unnecessary computational load.

Step size The maximum allowed step size between two successive frames defines the maximum distance that the MVA can pass through. The step size should be large enough to handle the possible motion, especially the high velocities in early diastole. However, it should not be too large to maintain temporal and spatial continuity of the motion [Üzümcü et al. 2006; van Stralen et al. 2005].

The values of the step size and the ROS size were chosen based on normal values of MVA displacement [Alam and Hoglund 1992; Behar et al. 2004]. These were consistent with the results of the manual tracking performed on the 20 patients for the evaluation of the proposed method. The chosen values are given in table 3.1.

Table 3.1: The applied values for the step size in X and Y direction (SS_x and SS_y) and the ROS

	SS_x (pixels)	SS_y (pixels)	ROS (pixel ²)
Septal	2	6	61×91
Lateral	2	7	61×91

3.2.4.1 | Parameter optimization

A systematic optimization procedure was carried out to select the values for the width of the smoothing filter and of the apodized kernel that lead to the best results in the sense that the Euclidean distance is minimized. The parameter optimization process was performed in the following fashion: the initial step was trying many possible combinations of the parameter values in a wide range and the three types of cost functions. From this phase, we extracted the type of cost function that leads to the best results and we could define a narrower range of the parameter values that should be involved in the systematic optimization. Based on the results of the initial step, all possible combinations of smoothing filter with a standard deviation in the range $[0, 3]$ and $[0, 2]$ pixels with steps of 0.5 pixels in the spatial and the temporal domain, respectively ('-' means that no smoothing at all was applied) and apodization kernel with a standard deviation in the range $[11, 19]$ pixels with steps of 2 pixels were examined. It is important to mention that the MDP search did not require any optimization procedure since it only depends on the ROS and the step size, which were determined beforehand.

Table 3.2: Optimal values of the parameters

	σ_x (pixels)	σ_y (pixels)	σ_t (frames)	α_x (pixels)	α_y (pixels)	ROS pixels ²
Septal	2.5	1	0.5	17	17	69 × 69
Lateral	-	-	-	17	13	69 × 53

σ_x , σ_y and σ_t are the standard deviations of the Gaussian filter in the spatial and the temporal domain, respectively, and α_x and α_y denote the standard deviations of the apodization kernel in the horizontal and vertical direction respectively. The region-of-interest (ROI) sizes are directly derived from α_x and α_y .

Data acquisition | 3.2.5

Single-beat apical four-chamber image sequences in 20 patients were recorded using a Hewlett-Packard Sonos 1500 or Sonos 2500 ultrasound system (Hewlett-Packard Inc, Andover, MA, USA) with a 2.5 MHz transducer. The patients were selected randomly from a database of 120 patients 3 days after acute myocardial infarction [Nijland et al. 2002]. Images were recorded on VHS videotape and digitized by a DT3152 frame grabber (Data Translation, Marlboro, MA, USA). The resolution was 768×576 pixels with pixel sizes ranging from 0.281×0.281 to 0.422×0.422 mm²/pixel. Cineloop images were acquired with the video frame rate of 25 Hz. All studies were acquired in apical position. To suppress high noise levels and interlacing artifacts in the raw data, a Gaussian filter (its parameters are given in table 3.2) was used to smooth the images in both the spatial and the temporal domain.

The tracking errors of three patients (for two of them in the lateral part and for the third one in the septal part) were detected as outliers and, therefore, they were excluded from this study. The failure of the tracking in two of them is attributable to the poor image quality of the sequences, resulting in a non discriminative ROI to be matched. As for the third outlier, after analyzing the results the two experts who performed the manual tracking agreed with the tracking results of the automated method rather than the averaged manual annotation. After excluding the three outliers from the study, the optimization was performed again with the remaining 17 patients (374 images in total).

outliers

3.2.6 | Data analysis

3.2.6.1 | Comparison with manual tracking

Since MVA tracking lacks “ground truth”, we compared the automatically computed positions with the positions indicated manually by two experts (JB and MvS). Both experts were constrained to begin the track from the same starting position and the rest of the cycle was annotated independently. The Euclidean distances between the computer generated annulus positions and the averaged manual positions were calculated. We used the average Euclidean distance as the distance measure.

3.2.6.2 | Comparison with other automatic tracking methods

The performance of the above automatic MDP-based procedure was evaluated relative to two other tracking methods:

1. Minimum tracking. This is the simplest approach for tracking problems based on block matching. The first step of this technique is applying block matching by calculating the costs using a single fixed template (ROI) for all frames (like in the proposed automatic procedure). After that, the optimal locations of the MVHPs are simply the local minima within the ROS for each frame, without enforcing any temporal continuity.
2. Forward tracking. Forward tracking is a common method that was used in previous studies [Behar et al. 2004; Eto et al. 2005; Golemati et al. 2003; Veronesi et al. 2006] for tissue tracking. In this algorithm block matching is applied and the minimum of the cost function is found in each frame separately. The template ROI_1 is updated in each frame and is set to be the content of ROI_t that minimized the cost function in the previous frame, as described in section 3.2.1. In this algorithm continuity is imposed by limiting the step size.

equal parameter settings The evaluation of the performance of these two methods was done using the same parameters as in the MDP-based method since they are independent of the tracking method. The parameters that were optimized are the width of the smoothing filter, which depends on the image quality and the size of the ROI, which is linked to the size of the anatomical structures to be tracked.

3.2.7 | Statistical analysis

Since we had manual tracking of two experts, we used the average of both of them as the best estimate of the true value. The difference between the tracking results

of the experts is a measure for the interobserver variability. Evaluating the performance of the proposed automated method by calculating the correlation between the manual tracking and the automated tracking is not enough in this context since high correlations do not necessarily mean that the two methods agree [Bland and Altman 1986]. Therefore, in addition to the linear regression analysis, we applied Bland-Altman analysis and measured the agreement between the MDP-based automated technique and the manual tracking. A paired t -test was applied to verify the significance of the bias. We compared between the interobserver variability and the errors of the automated method by examining the 95% confidence intervals of the limits of agreement. The 95% confidence intervals for the limits of agreement were calculated according to:

*confidence
interval*

$$CI = 2S \pm t \sqrt{\frac{3S^2}{n}} \quad (3.4)$$

where S denotes the standard deviation, n is the number of observations, and t is the 0.05 critical value of the t -distribution with $n - 1$ degrees of freedom. Since we compare to the average of multiple manual values, we would underestimate the limits of agreement with respect to the manual method. Therefore, a corrected standard deviation [Bland and Altman 1986] was used in the calculations of the limits of agreement between the automatic technique and the manual tracking:

*corrected
standard
deviation*

$$S_c = \sqrt{S_e^2 + \frac{1}{4} S_m^2} \quad (3.5)$$

where S_c is the corrected standard deviation, S_e is the original standard deviation of the automated errors and S_m is the standard deviation of the differences between the tracking results of the two experts.

Results | 3.3

Parameter settings | 3.3.1

The NCC as a matching criterion yielded the best results of tracking both the septal and the lateral parts and, therefore, it was chosen as the cost function to be applied. The parameters obtained with the parameter optimization process are shown in table 3.2.

NCC

3.3.2 | Tracking results

Figure 3.5 shows the average displacement over all 17 patients as function of the relative phase of the cardiac cycle. It can be seen that good agreement between the displacements found by the MDP-based automated method and the manual method was obtained.




Figure 3.5: The average displacements over all 17 patients in *Y* (left) and *X* (right) directions in the septal and lateral MVHPs as function of the relative phase of the cardiac cycle in percentage

Table 3.3: The maximal excursion as calculated by the automated method and the manual method

$n = 17$	MDP-based automated method (mm)	Manual method (mm)
Septal	11.0 $\pm$ 1.9	10.0 $\pm$ 1.8
Lateral	12.4 $\pm$ 2.8*	11.1 $\pm$ 3.2

All the results are expressed as mean $\pm$ standard deviation.

* Denotes an automated result which is significantly different ($p < 0.05$) from the manual result.

Maximal excursion | 3.3.3

The septal and the lateral MVHP maximal excursions over the cardiac cycle were calculated, the results are given in table 3.3. The MVA maximal excursion, which usually occurs at end systole, is the maximum distance between the initial MVHP (at end diastole) and the MVHP in all other frames. The first evaluation of the MDP-based automated method was done based on the maximal excursion results since it is related to other important clinical measures such as the ejection fraction and stroke volume [Alam and Hoglund 1992; DeCara et al. 2005; Eto et al. 2005; Willenheimer et al. 1999]. Bland-Altman plots of the difference between the calculated maximal excursions by both methods against the mean values of the methods are shown in fig. 3.6. In the septal part, no significant difference ($p > 0.05$) was found between the methods. As for the lateral part, a small but significant ($p < 0.05$) over-estimation of the automated method with a bias of 1.21 mm was found. However, it is important to mention that the maximal excursions in the lateral part, as indicated manually by both experts, were also found to be significantly different ($p < 0.05$) from each other.

Full cycle tracking | 3.3.4

The average MVA displacement errors (average Euclidean distance from the manual annotation) over all time samples and all 17 patients obtained from the proposed MDP-based automated procedure, from the forward tracking, from the minimum tracking and the interobserver variability are given in table 3.4. It can be clearly seen that the MDP-based tracking outperformed the forward and the minimum tracking.

High correlations were found ($p \ll 0.01$) between the displacement measure-




Figure 3.6: Bland-Altman plots of the difference between the maximal excursions in the septal (*left*) and the lateral (*right*) parts calculated by the automated method and those obtained by the manual method

Table 3.4: The errors obtained with the MDP-based automated method, the errors obtained from forward tracking, the errors obtained with the minimum tracking and the interobserver variability

<i>n</i> = 17	MDP-based tracking (mm)	Forward tracking (mm)	Minimum tracking (mm)	Interobserver variability (mm)
Septal	1.33 ± 0.86	2.14 ± 1.12	1.83 ± 2.19	1.43 ± 0.82
Lateral	1.66 ± 1.03	3.62 ± 2.21	4.21 ± 5.29	1.57 ± 0.94

All the results are expressed as mean ± standard deviation.

ments obtained by the manual method and by the MDP-based automated method. The results, including the regression lines and the correlation coefficients, are shown in fig. 3.7

Bland-Altman plots for the agreement between the displacements obtained by our method and by those obtained manually are given in fig. 3.8 and fig. 3.9. The displacement is defined as the offset from the position in the first frame (end diastole). In these plots a negative displacement in *Y* direction is an upward movement and a negative displacement in *X* direction describes movement to the left. In the lateral part (fig. 3.9), Bland-Altman analysis showed a small but significant underestimation (*p* < 0.05) with a bias of 0.33 mm in the displacements in *X* direction

coordinate
system




Figure 3.7: Results of the linear regression of the MVHPs displacement measurements obtained by the MDP-based automated method and by the manual method

and a small but significant overestimation with a bias of -0.52 mm in the displacements in Y direction. In the septal part (fig. 3.8) no significant difference ($p > 0.05$) was found.

The 95% confidence intervals of the limits of agreement, as calculated using eqn. 3.4, of both methods are shown in table 3.5. It can be seen that in the Y direction narrow limits of agreement for the difference between the two techniques were obtained although they are significantly different.

The computation time required for the proposed automatic method as well as those required for forward tracking and minimum tracking are given in table 3.6. The computation times shown are for full cycle tracking of the MVHPs on a notebook PC with an Intel Pentium M processor running at 1.73 GHz with 512 MB of RAM. However, we did not spend much effort in optimizing our method regarding

*computation
time*




Figure 3.8: Bland-Altman plots of the differences between the automated and the manual displacements in the septal part, in both the horizontal (X) and vertical (Y) direction




Figure 3.9: Bland-Altman plots of the differences between the automated and the manual displacements in the lateral part, in both the horizontal (X) and vertical (Y) direction

Table 3.5: The 95% confidence intervals for the limits of agreement of the errors in X direction and Y direction, respectively

<i>n</i> = 17	MDP-based tracking		Interobserver variability	
	X (mm)	Y (mm)	X (mm)	Y (mm)
Septal	2.28 ± 0.20	1.86 ± 0.16*	2.39 ± 0.21	1.63 ± 0.14
Lateral	3.12 ± 0.27	2.55 ± 0.22*	2.92 ± 0.26	2.09 ± 0.18

All the results are expressed as mean ± standard deviation.
*Denotes an automated limit of agreement which is significantly different (*p* < 0.05) from the interobserver variability.

Table 3.6: Computation times per patient required for the three tracking methods

<i>n</i> = 17	MDP-based tracking (s)	Forward tracking (s)	Minimum tracking (s)
Computation time	161.1 ± 10.3	110.1 ± 9.7	143.7 ± 13.5

All the results are expressed as mean ± standard deviation.

execution speed.

Discussion | 3.4

We developed a semi-automatic technique for tracking the MVA motion using MDP combined with apodized block matching. We applied this method on clinical quality VCR recorded data from 20 patients. We evaluated its feasibility on 17 patients, after excluding three patients due to poor image quality or ambiguity of the manual annotation of either the septal or the lateral MVHP. Tracking the MVA based on manual annotation is a cumbersome and tedious procedure with poor repeatability. To overcome this limitation, we propose a semi-automatic method that requires minimal user interaction. Relative to forward tracking and minimum tracking, the method is robust to noise and small disturbances.

Our method has several advantages for evaluating the MVA motion. First, it exploits the DP advantages: relatively insensitive to noise and small disturbances, efficient and producing a time continuous motion [Rabben et al. 2000; Sonka et al. 1999]. Another benefit from the MDP search is the 2D path derived from it. In this way, our technique is capable of evaluating the spatial displacement of the MVA in both the horizontal and the vertical direction over the full cardiac cycle. Second, it takes the advantages of block matching: robust against linear changes, fast and proven. Third, it is simple and can be performed off-line with a PC. In addition, the method has the benefits of 2D echocardiography: cheap relative to other image acquisition modalities, non-invasive, accessible and the images are produced in real time.

advantages

Previous studies [Alam and Hoglund 1992; DeCara et al. 2005; Eto et al. 2005; Willenheimer et al. 1999] reported that the MVA maximal excursion correlates well with LV function. The maximal excursions values obtained with our automated method agreed well with the values obtained with the manual method, consider-

MVA excursion

ing the inherent uncertainty of the manual drawing. Eto et al. [2005] developed an automated method, based on forward tracking, for tracking the MVA in high frame rate (greater than 60 Hz) cineloop images of 2D echocardiography. They validated their method by comparing its results to those obtained by manual annotation in 3D echocardiography and reported an accuracy of 2.5 ± 1.8 mm (the average distance between the mean MVA excursion obtained by the automated method and the manual tracking). However, it is difficult to compare these results to those obtained by our method since they used a different “ground truth”.

*full cycle
tracking results*

The full cycle tracking errors obtained with the MDP-based automated technique were small (less than 2 mm, see table 3.4) and with clinically acceptable differences [Eto et al. 2005]. The tracking errors in the lateral part were larger than in the septal part. This can be explained by the fact that the lateral part is more difficult for tracking since the ROI has less clearly defined structures with lower echo level. Consequently, the interobserver variability was also larger in the lateral part than in the septal part. We found limits of agreement of less than 2 mm which is good given the uncertainty in our “ground truth”, which is expressed by the manual interobserver variability. Bland-Altman plots revealed good agreement between the displacements in X direction obtained by both methods. However, the limits of agreement of the displacements in Y direction obtained by the automated method were narrow but significantly different from those obtained by the manual method (see table 3.5). The wider limits of agreement in the X direction are attributable to the resolution, which is higher in the axial direction than in the lateral direction. Therefore, the resolution in the X direction, which is mainly governed by its lateral component, is worse and the limits of agreement are broader.

The displacement in the Y direction is clinically more valuable than in the X direction. No apparent correlation between the error and the average displacement was found (fig. 3.8 and fig. 3.9). This is a great advantage of the MDP-based automated method, which does not suffer from increasing errors even at large excursions.

The fact that the best results in the lateral part were obtained without smoothing the images implies again that in the lateral part there are no clear anatomical structures for tracking and the MVHP is tracked based on the speckles.

*related
techniques*

In our study, the MDP-based automated method outperformed the forward tracking and the minimum tracking while requiring comparable computation time (table 3.6). The forward tracking approach is fastest since the ROS is effectively smaller (limited by the step size), at the cost of finding a suboptimal solution. Unlike other echocardiography studies [Behar et al. 2004; Eto et al. 2005; Rabben et al. 2000], we used cineloop images with low frame rate (25 Hz) compared with digital images. Forward tracking was reported to perform well in high quality images [Behar et al. 2004; Eto et al. 2005; Veronesi et al. 2006] but yielded poor results in

our case. This is attributable to the fact that in forward tracking once an error occurs it is emphasized since the new ROI is updated erroneously and the detected path drifts away from the true path. An important contribution to this problem is the change in appearance of the mitral valve leaflet due to opening and closing of the valve. Since this is the nature of the problems, applying our method on digital images with higher frame rate and better image quality would probably reduce the tracking errors. This is attributable to the fact that our method searches for globally optimal location of the MVHPs using MDP and uses apodized block matching to eliminate the dominant contribution from peripheral pixels in the ROI.

We chose the step size and the ROS size based on normal ranges of MVA displacement [Alam and Hoglund 1992; Behar et al. 2004] and, hence, they do not depend on the quality of the images used in this study. The performance of the method was not very sensitive to the choice of the parameters (smoothing filter size and ROI size) that were selected during the optimization procedure; the differences were only a few percent. However, these parameters depend on the image quality and hence we expect to find different values for other types of data, e.g. high-quality digital images.

*choice of
parameters*

We used the same parameters for tracking the entire data set. Better results may be achieved by selecting the ROI size for each patient separately based on the autocorrelation function of the system [Behar et al. 2004]. When implementing the proposed method on high quality digital images, its results may be further improved by using a nonlinear edge preserving filter to increase the signal to noise ratio (SNR) of the images while preserving important image features [Adam et al. 2006; Behar et al. 2004] instead of smoothing the data with a linear Gaussian filter as we did. Using nonlinear filters would probably improve the results in the septal part, where the MVHP is a strong reflector, rather than in the lateral part, where the MVHP is weaker (Behar et al. 2004). In our images with dominant interlacing artifacts, however, a smoothing filter is better than an edge preserving one.

*further
improvements*

Further work can be the extension of the method to 3D data. The MDP can be easily adapted to track the MVA in any long-axis slice of a four-dimensional data set (3D + time), or for detecting the complete mitral ring in a 3D set.

Limitations | 3.4.1

We analyzed only data obtained from patients with acute myocardial infarction. Although this subsample might not be fully representative, we do not see any reason why this method would not be suitable for normals or for patients with other pathologies, as long as no extreme kinematical changes are involved.

*patient
population*

The quality of the images was relatively low due to the VCR recorded data. Hence,

image quality

this study should only be seen as evidence for the feasibility of MDP-based automatic mitral valve tracking in VCR images. However, we developed this method on such relatively poor-quality data because we were aiming for a robust technique suitable for 3D ultrasound data, with an inherently lower quality than modern digital data. Since we wanted to use gross anatomical features rather than highly correlated speckle patterns in high frame rate, digital images with good SNR, the VCR images offered a good substrate.

The suitability of this technique for 3D echo still needs to be investigated, but its robustness is promising. We believe that our approach will also be useful on digital data, but further work on such data should be undertaken. At least, it will require a new parameter optimization process since the image quality influences the smoothing filter parameters.

3.5 | Conclusions

We developed a novel semi-automatic method based on MDP for tracking the MVA motion. This study showed that the technique is feasible and valuable in evaluating accurate MVA displacements. The method provides robust and efficient tracking with minimal user interaction and agrees well with manual MVHPs tracking even in low quality images.

| Acknowledgments

The authors acknowledge Dr. F. Nijland and Dr. O. Kamp (Department of Cardiology, VUMC, Amsterdam, the Netherlands) for supplying the data that was used in this study.

Interpolation of irregularly distributed sparse 4D ultrasound data using normalized convolution

WE DEVELOPED A NOVEL MULTI-BEAT image fusion technique using a special spatiotemporal interpolation for sparse, irregularly sampled data (ISIS). It is applied to irregularly distributed 3D cardiac ultrasound data acquired with a fast rotating ultrasound (FRU) transducer. ISIS is based on normalized convolution with Gaussian kernels tuned to irregular beam data spacing over cardiac phase (τ), and beam rotation (θ) and elevation angles (ϕ).

Images are acquired with the FRU transducer, a linear phased array rotating mechanically continuously at very high speed (240-480rpm). High-quality 2D images are acquired at ± 100 frames/s over 5-10 seconds. ECG is recorded simultaneously. Images are irregularly distributed over τ and θ , because rotation is not synchronized to heart rate. ISIS was compared quantitatively to spatiotemporal nearest neighbor interpolation (NNI) on synthetic (distance function) data of a pulsating ellipsoid. ISIS was also tested qualitatively on 20 in-vivo cardiac image sets and compared to classical temporal binning with trilinear voxel interpolation, at resolutions of $256 \times 256 \times 400$ for 16 phases.

From the synthetic data simulations, ISIS showed absolute distance errors (mean $\pm$ standard deviation) of 1.23 ± 1.52 mm; considerably lower than for NNI (3.45 ± 3.03 mm). For in-vivo images, ISIS voxel sets showed reduced motion artifacts, suppression of noise and interpolation artifacts and better delineation of endocardium. ISIS improves the quality of 3D+T images acquired with a FRU transducer in simulated and in-vivo data.

This chapter has been derived from (© 2006 SPIE):

Novel spatiotemporal voxel interpolation with multibeat fusion for 3D echocardiography with irregular data distribution. J.G. Bosch, M. van Stralen, M.M. Voormolen, B.J. Krenning, C.T. Lancé, J.H.C. Reiber, A.F.W. van der Steen, N. de Jong. Proc SPIE Med Imaging 2006; 6147: 61470Q.

4.1 | Introduction

3D ultrasound techniques

3D ultrasound imaging is rapidly gaining clinical importance. For acquisition of 3D images, several techniques are possible: freehand probe manipulation with position sensing; motorized probe translation/rotation; or matrix array transducers. For acquisition of 3D sets over the full cardiac cycle (4D echocardiography), this is generally combined with ECG triggering. The speed and quality of a freehand or stepper motor approach are insufficient for such applications. As an alternative to the costly and complex matrix array solution, we have developed a linear phased array transducer that rotates continuously at high speed. Such a fast rotating linear phased array can be used to acquire images of the left ventricle covering 3D space over the full cardiac cycle in near real time. Image quality within the 2D image planes is superior to matrix array transducers, but these planes are spread sparsely and irregularly over space and time, since rotation and cardiac phase are not synchronized. We developed a novel spatiotemporal interpolation technique to tackle the complex problem of generation of optimal voxel sets and standardized cross sections from this data.

the challenge

4.1.1 | Fast rotating ultrasound transducer

transducer design

The fast rotating ultrasound (FRU) transducer developed in our laboratory [Djoa et al. 2000; Voormolen et al. 2002; Voormolen et al. 2003] is a special device for near real-time 3D cardiac image acquisition (fig. 4.1). It consists of three major parts: a phased array, a DC motor that drives the array, and a slip-ring device establishing signal transfer to and from the continuously rotating array. The DC motor drives the array at a high rotation speed ranging from 4 to 8 Hz and is connected to an external control system with a manual setting for the rotation speed. The array of the transducer, custom made by Delft Instruments (Delft, The Netherlands), contains 64 elements with a pitch of 0.21 mm and is tapered into an octagonal shape, approximating a circle with a radius of 7 mm. It has a fractional bandwidth of 86% with a center frequency of 3.2 MHz.

4.1.2 | Image acquisition

Cardiac apical images are acquired at about 100 2D frames per second (fps) with the FRU transducer rotating at 4–8 Hz (240–480 rpm) and recorded with a GE/VingMed Vivid 5 ultrasound system (VingMed, Horten, Norway). The typical acquisition



Figure 4.1: The fast rotating ultrasound (FRU) transducer

time is approximately 10 seconds, which has proven to be convenient for clinical application and allows a complete acquisition within one breath hold.

2D images are acquired and stored digitally in polar format (ultrasound beams). Each image consists of approximately 80-90 beams at about 1° spacing; each beam consists of 400-500 samples. The patient's ECG is recorded simultaneously. Due to the fast rotation, individual 2D images have a curved 3D shape (fig. 4.2) because the image plane rotates over $10^\circ - 30^\circ$ during the collection of the consecutive beams of the image.

*curved image
planes*




Figure 4.2: A sequence of seven consecutive FRU images with curved image planes

Because of variable cardiac cycle duration and lack of synchronization between rotation and cardiac phase, images are irregularly distributed over rotation angle and cardiac phase. After acquisition, exact rotation angle (θ) and cardiac phase (τ , time offset from R-peak) is calculated for each 2D image frame and its beams. From this information, multi-beat fusion can be applied: around 10 consecutive cardiac cycles are merged into one to reach a denser coverage of the 4D space.

4.1.3 | Sparse data and multi-beat fusion

In fig. 4.3a, an example of the coverage of 3D space is shown for an interval of about $1/12^{\text{th}}$ of a cardiac cycle. It can be seen that in that period, the volume is swept almost two times. At 6 rotations per second and 100 fps, a full volume is swept in about 80 ms and contains about 8 images (16 half sectors). The 3D distances between images amount to 10 – 20 mm, which is too large for appropriate voxel space generation; therefore, corresponding cardiac phases from several cardiac cycles are merged. This multi-beat fusion results in a denser coverage of 3D space per cardiac interval (fig. 4.3b).




Figure 4.3: *left*) Example of the sparse sampling of the time-volume space resulting from one time interval ($1/12$) of a cardiac cycle. The individual recorded images are depicted as lines at a certain C-plane level (the z -axis is the transducer's axial direction). The left half of each individual sector image is shown as a solid line and the right half as a dashed line. *right*) The denser sampling of the volume-time space when multi-beat fusion is used

Irregular data distribution | 4.1.4

The distribution of the curved image planes over the 4D data space (3D plus time) is determined by three factors: the rotation speed, the frame rate and the duration of each cardiac cycle. The frame rate is determined by the settings of the ultrasound machine, remains fixed over an acquisition and is exactly known. Rotation speed is set at the motor controller of the FRU transducer and measured precisely after the acquisition by a special image analysis procedure [Voormolen et al. 2003]. The beat lengths, however, are variable and unpredictable, since they are determined by the patient's spontaneous sinus rhythm, which exhibits a natural short-term variability. Therefore, the start of the next beat (marked by the R-peak in the ECG) comes at an unpredictable moment, and thus at a highly unpredictable rotation angle. Consecutive beats may easily vary in beat length by 50 ms, which corresponds to a rotation angle difference of about 100° . An actual example of the irregular distribution of images over rotation angle θ and cardiac phase τ is given in fig. 4.4.

*unpredictable
beat length*

Interpolation of sparse, irregular samples | 4.1.5

Interpolating data from sparse and irregularly distributed samples is a classical problem. In the image processing domain, it has been addressed for the filling of missing parts in images with “plausible” data (image inpainting) [Drori et al. 2003; Pham and van Vliet 2003], resolution improvement by combining multiple low-resolution images into one high-resolution images (superresolution) [Farsiu et al. 2004; Pham and van Vliet 2003] and converting irregularly sampled data into regular images [Loke and du Buf 2004]. In the (ultra)sonic imaging domain, such issues have been addressed for sonar imaging [Loke and du Buf 2004] and interpolation of freehand 3D ultrasound [Estépar et al. 2003; Meairs et al. 2000]. For inpainting and superresolution applications, the so-called normalized convolution [Knutsson and Westin 1993; Westin 1994] approach (also known as normalized averaging) has been applied with success [Drori et al. 2003; Pham and van Vliet 2003].

*normalized
convolution*

The technique has been applied to freehand 3D ultrasound images of cerebrovascular structures [Meairs et al. 2000] and kidney and thyroid [Estépar et al. 2003] as well, but these images were relatively densely sampled and there was no temporal interpolation involved. In our case, we deal with three spatial dimensions and one temporal dimension. The default technique in this 4D domain is temporal binning (assigning all data to some discrete time points) in combination with nearest neighbor interpolation or trilinear spatial averaging. To our knowledge, normalized convolution has not been applied to this 4D problem domain.

*application to
4D domain*




Figure 4.4: Distribution of 1012 consecutive image frames (14 beats) over rotation angle and cardiac phase. Beat durations from 0.8 to 0.92 s, rotation speed 5.82 Hz, frame rate 106.1 Hz. Each dot represents one frame. The redlines highlight the consecutive frames from the first beat (from left to right); the light redline marks those of the second beat. The redcircles show areas of dense and sparse distribution

4.2 | Methods

4.2.1 | Introduction

We developed a novel multi-beat fusion technique using a special spatiotemporal interpolation for sparse, irregularly sampled data (ISIS). The technique is based on normalized convolution with Gaussian kernels tuned to the (irregular) spacing of the data over cardiac phase (τ) and over the rotation (θ) and elevation angle (ϕ) of the spherical beam coordinates. Normalized convolution is a technique known from superresolution approaches and irregular data interpolation, but has not been applied to this area. The interpolation of the sparse irregular beam data takes place in the (θ, ϕ, τ) domain; results are then converted into Cartesian (x, y, z) voxel space for each τ .

Image fusion/interpolation | 4.2.2

For constructing a 4D image set (a time sequence of 3D voxel sets), this irregularly and locally sparsely populated dataset must be interpolated into a regular, fully populated set. In fig. 4.4, this corresponds to estimating the data at each of the grid line crossings (or even on a finer grid). Various techniques can be applied to fulfill this task (fig. 4.5).




Figure 4.5: Different approaches to sparse data interpolation, demonstrated on a detail of fig. 4.4. *a)* grid position to be filled. *b)* Spatiotemporal nearest neighbor interpolation. *c)* temporal binning and spatial linear interpolation. *d)* spatiotemporal weighted average

The simplest approach is to find the closest sample for each grid point. This means we bluntly neglect the error in time and angle, which can result in considerable motion and fusion artifacts. Moreover, many of the available samples are simply neglected, since only one sample per grid point is used for the estimation. Furthermore, several grid points may use the same single sample. In terms of interpolation, this is a spatiotemporal “nearest neighbor” approach.

*nearest
neighbor*

The interpolation in our case involves both temporal and spatial differences. A sensible classical solution is to apply temporal binning followed by interpolation. All samples within a time interval are assigned to the midpoint of the interval, but their angular information is preserved. Interpolation is then applied between consecutive angles. Although no samples are wasted and angular information is preserved, this may still result in considerable motion artifacts, since information from different time points is combined without taking object motion into account.

*temporal
binning*

An alternative solution for spatiotemporal interpolation is to use a weighted combination of all neighboring samples in space and time. For sparse, irregularly distributed data, finding the neighboring samples for each grid point is not an easy task. Therefore, most approaches use the inverse approach and distribute the sample data to neighboring grid points.

our approach We have developed a novel technique using a dedicated spatiotemporal interpolation for sparse, irregularly sampled data (ISIS). The technique is based on normalized convolution (NC) [Knutsson and Westin 1993; Pham and van Vliet 2003; Westin 1994]. Normalized convolution applies the signal/certainty principle: a sparse data sampling can be represented as a combination of a signal image S , containing the available samples, and a certainty image C , containing value '1' for the sample positions and value '0' elsewhere.

The empty spaces can be filled by diffusing both signal strength and certainty around the available samples. By normalizing the accumulated samples with accumulated certainty, properly weighted values result. NC is a technique known from superresolution approaches [Knutsson and Westin 1993; Pham and van Vliet 2003] and irregular data interpolation [Knutsson and Westin 1993], but has not been applied to this area of spatiotemporal ultrasound interpolation.

2D example In the example of fig. 4.6, we will use the distribution shown in fig. 4.4. Each sample point in this distribution represents one complete image (all values of r and ϕ). Therefore, with this example we can simulate the interpolation over the (θ, τ) domain which is sparsely and irregularly sampled. By applying the distribution of the 1012 samples on the 128×128 test image (fig. 4.6a), we keep only 6% of the original data (fig. 4.6b and c). By applying a convolution ($\otimes$) with some diffusion kernel G on both the signals (fig. 4.6d) and their certainty (fig. 4.6e) and calculating their ratio (eqn. 4.1, fig. 4.6f), we get the normalized convolution result R .

$$R = \frac{(S \cdot C) \otimes G}{C \otimes G} \quad (4.1)$$

In fig. 4.6, a 2D Gaussian kernel was applied with a width $\sigma = 3$ in both directions. As can be seen in fig. 4.6f, the structure of the image is well preserved, although only 6% of the original samples were used.

4.2.3 | Implementation in 4D

Instead of the 2D example described above, we are addressing a 4D problem: one continuous temporal dimension and three spatial dimensions. In our case, the goal of the interpolation is the generation of a time sequence of regular Cartesian (x, y, z) voxel sets. The input data, however, consists of discrete ultrasound beams; their native domain is the spherical coordinate system, where every beam is fully characterized by its rotation angle θ , elevation angle ϕ and cardiac phase τ . Each sample within the beam is addressed by its depth r (fig. 4.7a).

We propose to apply spatiotemporal interpolation between complete beams, so the variable r can be left out of the interpolation task. The interpolation of the




Figure 4.6: Example of normalized convolution. *a)* original image 128×128 . *b)* sparse sampling S of image following distribution of fig. 4.4. *c)* Certainty image C . *d)* image S diffused by Gaussian kernel. *e)* image C diffused by Gaussian kernel. *f)* Result image, D normalized by E




Figure 4.7: *a)* Spherical and Cartesian coordinate system. The z -axis corresponds to the rotation axis of the transducer. r : depth, ϕ : elevation angle, θ : rotation angle. A single curved image plane with a scan sector and a single beam is shown in gray. A spherical plane $r = C$ is shown in dark gray. *b)* Domain of ISIS interpolation: for planes of constant r , ISIS is applied over (θ, ϕ, τ) space

sparse beam data takes place in the (θ, ϕ, τ) domain (fig. 4.7b), by 3D normalized convolution using a 3D Gaussian kernel. This results in a dense, oversampled, regular (θ, ϕ, τ) space, or a regular, dense (θ, ϕ) plane for each desired value of τ . Results are then converted into Cartesian (x, y, z) voxel space for each τ by trilinear interpolation in a regular grid between the eight closest points in (θ, ϕ, r) for every voxel (fig. 4.8).




Figure 4.8: Conversion from spherical to Cartesian domain: ISIS interpolated sets for all r are combined and converted into regular voxel sets for each desired value of τ

4.3 | Experiments

ISIS was tested both quantitatively (by simulations) and qualitatively on in-vivo data. It was tested quantitatively by simulations on synthetic data representing the beating left ventricle and compared to spatiotemporal nearest neighbor interpolation (NNI). From the synthetic data, spatiotemporal sampling errors D of both methods were calculated by sampling a distance function from the beating ellipsoid's surface and subtracting the true (known) distance at the exact desired time and spatial position. ISIS voxel space generation was also tested qualitatively on 20 in-vivo cardiac image sets of 15 patients, compared to a classical temporal binning with spatial trilinear voxel interpolation.




Figure 4.9: Results of simulations on a distance function for a pulsating ellipsoid over (θ, τ) space. *a*) Distribution of samples (beams) over phase τ and rotation angle θ (300×256). *b*) Distance function sampled at ideal positions for θ and τ (37×32). *c*) NNI results (37×32) from samples in *a*. *d*) ISIS results (37×32) from samples in *a*

Quantitative evaluation on synthetic data | 4.3.1

ISIS was tested quantitatively by simulations on synthetic data of a pulsating ellipsoid representing the beating left ventricle. It was compared to spatiotemporal nearest neighbor interpolation (NNI). For every point $\mathbf{p}$ in space at time t , an analytical distance function $D(\mathbf{p}, t)$ from the ellipsoid's surface can be calculated,

error measure

$$D(\mathbf{p}, t) = \frac{(\|\mathbf{p} - \mathbf{a}\| + \|\mathbf{p} - \mathbf{b}\|)}{2} - D_0(t) \quad (4.2)$$

where $\mathbf{a}$ and $\mathbf{b}$ are the 3D positions of the ellipsoid's foci and $D_0(t)$ is the ellipsoid's long-axis radius as a function of time.

From the synthetic data, spatiotemporal sampling errors ΔD of both interpolation methods were calculated by sparsely sampling the distance function from the beating ellipsoid's surface (eqn. 4.2), interpolating the sparse distance data and comparing the interpolated distance values $D'_{NNI}(\mathbf{p}, t)$ resp. $D'_{ISIS}(\mathbf{p}, t)$ with the true (known) distance $D(\mathbf{p}, t)$ at all exact desired spatiotemporal positions $(\mathbf{p}, t)$,

$$\Delta D = D'(\mathbf{p}, t) - D(\mathbf{p}, t) \quad (4.3)$$

ΔD expresses both temporal and spatial errors in a distance in mm, which is meaningful since the beating ellipsoid's dimensions model the actual application.

Target resolutions for the rotational image sets were 32 angles (θ) and 37 phases (τ) within the cardiac cycle (which corresponds to resolutions of about 11° , 25 ms) (fig. 4.9). Object displacements per time step were in a similar range as per angle step. For simplicity, this evaluation was only done in two dimensions (one temporal, one spatial); the images can be interpreted as an M-mode of a set of 32 samples on a coaxial circle in a C-plane (horizontal; ϕ, r fixed; θ varying).

4.3.2 | Qualitative evaluation on in-vivo data

ISIS voxel space generation was also tested qualitatively on 20 in-vivo cardiac image sets of 15 patients, compared to a classical temporal binning with spatial trilinear voxel interpolation. Target spherical resolutions for normalized convolution interpolation over (θ, ϕ, τ) were $(256 \times 100 \times 64)$, in order to achieve (x, y, z) voxel sets of $(256 \times 256 \times 400)$ at 16 cardiac phases. Typical dimensions of the input data were 800 – 1100 polar frames containing 80 – 90 beams with 400 – 500 samples, covering a sector of about 90° and 10 cm depth. The applied 3D Gaussian diffusion kernel $G(\sigma_\theta, \sigma_\phi, \sigma_\tau)$ had widths of (1.0, 1.0, 3.0) respectively. A qualitative evaluation was performed by side-by-side comparison of equivalent cross sections (fig. 4.10) of voxel sets generated by both techniques.

parameter
settings




Figure 4.10: *top*) Reconstructed planes using classical trilinear interpolation in 3D space, assuming constant cardiac phase within each time bin. *bottom*) same planes using ISIS. Depth ± 10 cm, just above mitral valve at (left to right) phase ED, ES and iso-volumetric contraction. ISIS gives higher tissue-to-blood ratio and a decrease in motion artifacts

Table 4.1: Results of interpolations on synthetic data of pulsating ellipsoid

$n = 17$	Distance errors (mm)	Absolute distance errors (mm)
NNI	0.31 ± 4.58	3.45 ± 3.03
ISIS	0.16 ± 1.95	1.23 ± 1.52

All the results are expressed as mean $\pm$ standard deviation.

Results | 4.4

Simulations on synthetic data | 4.4.1

In the simulations, nearest neighbor interpolation (NNI) showed temporal errors of -1.1 ± 13.8 ms (mean $\pm$ standard deviation) and angular errors of -0.2 ± 4.8 degrees. Mean and SD of absolute errors were 10.2 ± 9.4 ms and 3.8 ± 3.0 degrees, respectively. These values are directly calculated from the nearest neighbor sparse sample positions. From the synthetic data simulations, ISIS showed distance errors ΔD_{ISIS} of 0.16 ± 1.95 mm; this was considerably lower than for NNI ($\Delta D_{NNI} = 0.31 \pm 4.58$ mm). Absolute distance errors were 1.23 ± 1.52 mm for ISIS vs. 3.45 ± 3.03 mm for NNI. Results are shown in fig. 4.9 and listed in table 4.1.

In-vivo images | 4.4.2

In the in-vivo images, the resulting voxel sets showed reduced motion artifacts, suppression of noise and interpolation artifacts and better delineation of structures and tissue-to-blood contrast (fig. 4.10). Especially the typical zigzag artifacts (resulting from interpolations between frames that are spatially very close but temporally different) are strongly suppressed. However, the choice of the kernel size parameters seemed suboptimal in some cases. In areas of very sparse sampling, some artifacts remain visible. Sometimes, interpolation seems irregular in these areas and spurious edges may result. In very densely sampled areas, sometimes some anatomical details are smoothed out.

4.5 | Discussion

As the most important result from this study, we can conclude that the ISIS method gives better interpolation results than a classical approach using temporal binning. However, several limitations apply and we assume further improvements are possible; these are discussed below.

4.5.1 | Limitations

The setup of this study was limited in several aspects.

simulation only in 2D First of all, the simulations have so far been applied only in the 2D case. Although we have chosen this simulation in a way that matches the real problem well, we would like to extend it to the actual 4D situation.

qualitative in-vivo evaluation Secondly, the in-vivo interpolation results have so far only been judged qualitatively by visual comparison to other interpolations. A more elaborate evaluation needs to be performed, e.g. on a moving phantom, to objectively investigate artifacts and to examine the effects of the method on volume estimations after segmentation.

4.5.2 | Alternatives and further development

There are several issues that should be further investigated.

size of the diffusion kernel As described in section 4.4.2, the sizes of the Gaussian diffusion kernels seem suboptimal in some situations. The size of the diffusion kernel should match the local sample distribution; if the kernel size is much smaller than the spatiotemporal sample distances, the diffusion process can result in an irregular interpolation, possibly introducing spurious edges. If the kernel size is much larger than these local sample distances, data from many points will contribute to fill the holes and results will look overly smoothed. This can be understood if we consider the weighted sum of two overlapping Gaussian kernels at different distances. Where sample points are relatively close, the Gaussian kernels will result in an equal weighting and therefore in an unweighted averaging; at the tails of the distributions, the nearest sample will fully dominate the weighting and the result will behave like a nearest neighbor interpolation.

locally adaptive kernel This issue can be addressed by applying alternative convolution kernels, such as proposed by Knutsson and Westin [1993] and Westin [1994], or by locally adjusting the kernel sizes to the distribution densities, e.g. by using a distance transform to

determine the widths of the anisotropic diffusion kernels. Preliminary efforts in this direction showed promising results.

The ISIS method fulfills its tasks in the sense that it supplies plausible spatiotemporal interpolation of sparse, irregular data. It effectively reduces noise and suppresses motion artifacts. However, a technique like this will not always preserve sharp edges. Especially in regions of sparse sampling and/or considerable motion, edges will inevitably be smoothed. Spatiotemporal interpolation of a moving, sharp edge will result in a smoothed or double edge. This could be handled with specific assumptions about objects and edges and a certain motion model, but this is subject of further research. Also, registration could be combined with the interpolation (e.g. as proposed by Penney et al. [2004]) to address this limitation. However, such an approach is iterative, computation intensive and known to be extremely slow. Therefore, this will result in considerable further increase in complexity and processing time.

Conclusions | 4.6

The proposed interpolation technique (ISIS) improves the quality of 3D + T images acquired with a fast rotating transducer in simulated and in-vivo data. Several further improvements seem feasible. The technique may also prove useful in similar cases where multi-beat fusion, irregular sampling and/or phasic motion artifacts induce image errors, e.g. freehand 3D ultrasound.

Automatic time continuous detection of the left ventricular long axis and the mitral valve plane in 3D echocardiography

AUTOMATED SEGMENTATION APPROACHES for the left ventricle (LV) in 3D echocardiography (3DE) often rely on manual initialization. So far, little effort has been put in automating the initialization procedure to get to a fully automatic segmentation approach.

We propose a fully automatic method for the detection of the LV long axis (LAX) and the mitral valve plane (MVP) over the full cardiac cycle, for the initialization of segmentation algorithms in 3DE. Our method exploits the cyclic motion of the LV and therefore detects salient structures in a time-continuous way. Probabilities to candidate LV center points are assigned through a Hough transform for circles. The LV LAX is detected by combining dynamic programming detections on these probabilities in 3D and 2D + time to obtain a time continuous solution. Subsequently, the mitral valve plane is detected using the previously detected LAX.

Automatic detection was evaluated using patient data acquired with the fast rotating ultrasound (FRU) transducer (11 patients) and with the Philips Sonos 7500 ultrasound system, with the X4 xMatrix transducer (14 patients). For the FRU data, the LAX was estimated with a distance error of 2.85 ± 1.70 mm (mean $\pm$ standard deviation) and an angle of 5.25 ± 3.17 degrees; the MVP was estimated with a distance of -1.54 ± 4.31 mm. For the matrix data, these distances were 1.96 ± 1.30 mm with an angle error of 5.95 ± 2.11 and -1.66 ± 5.27 mm for the mitral valve plane. These results confirm reliable detection of the LV LAX and MVP. It may therefore serve as a replacement of manual initialization of 3D segmentation approaches.

This chapter has been derived from (© 2008, with permission from Elsevier):

Time continuous detection of the left ventricular long axis and the mitral valve plane in three-dimensional echocardiography. M. van Stralen, K.Y.E. Leung, M.M. Voormolen, N. de Jong, A.E.W. van der Steen, J.H.C. Reiber, J.G. Bosch. *Ultrasound Med Biol* 2008; 34(2): 196-207.

5.1 | Introduction and literature

Three-dimensional (3D) echocardiography is an increasingly widely available acquisition technique for assessment of left ventricular (LV) function, which is non-invasive, relatively cheap and portable. Due to the rapid increase of the use of this modality for diagnosing global and regional LV function, valvular disease, etc. there is a growing demand for objective, reproducible and automated techniques for identification of salient structures and quantification of left ventricular function.

Recently, several approaches for automated endocardial contour detection have been proposed, reporting success in measuring global functional parameters [Angelini et al. 2001; Corsi et al. 2002; Gérard et al. 2002; Köhl et al. 2004; van Stralen et al. 2005; Zagrodsky et al. 2005], and significantly decreasing the amount of user interaction that is needed for these measurements. Also, several quantification tools have already become available on commercial 3D echocardiography (3DE) systems. Although time can be gained by using these tools, reliable quantification of important clinical parameters is still very labor intensive and not yet ready for use in daily clinical routine. This requires techniques that need minimal or no user interaction. Moreover, automating the initialization procedure of such methods would also eliminate inter- and intraobserver variability from the analysis, increasing the reproducibility of the analysis.

Most previously presented methods for the quantification of LV function require some manual initialization. Initialization is done either by explicitly annotating the apex, a number of points on the endocardial border or the mitral valve [Corsi et al. 2002; van Stralen et al. 2005], or by indicating the LV position and dimensions by annotations [Angelini et al. 2001; Gérard et al. 2002; Köhl et al. 2004].

Although much attention has been paid to minimize this user interaction for automated contour detection, little effort has been put in developing dedicated automatic initialization procedures, which focus on automatically detecting salient structures in 3D echocardiography (3DE). Stetten and Pizer [1999] attempt to detect the apex and mitral valve center using medial-node models. Veronesi et al. [2006] describe a method for automated detection of the LV LAX based on optical flow, but it still needs manual initialization. The segmentation method by Zagrodsky et al. [2005] is initialized using a time-consuming registration with a pre-segmented template image.

Automated initialization has received more attention for segmentation of cardiac magnetic resonance (MR) images. We were inspired by work of van der Geest et al. [1997] and Müller et al. [2005], who detect the LV center in short-axis MR images using a Hough transform [Ballard 1981] for circles, for initialization of automated

endocardial border detection.

We propose a fully automatic method for reliable estimation of the position of the mitral valve and the orientation of the left ventricular long axis for apical 3D echocardiographic images of clinical quality. It is computationally inexpensive and easily adaptable, which makes it a valuable starting point for various high-level segmentation techniques.

*proposed
method*

Materials and methods | 5.2

We propose a method for finding the long axis (LAX) of the left ventricle and the mitral valve plane (MVP). Such a method should be capable of dealing with typical ultrasound acquisition characteristics, such as inhomogeneous image intensities, speckle and partial dropout of the myocardium.

The detection of these salient structures of the LV is achieved by a few robust consecutive steps (fig. 5.1). At first, the LAX is detected by locating the main circular structure in a number of planes perpendicular to the (apical) acquisition axis over time, using a Hough transform for circles (fig. 5.1a-b). Consecutively, multidimensional dynamic programming is applied in 3D and 2D + time to locate probable LV centers (fig. 5.1c-e). Fitting a line through the LV centers in each cardiac phase results in the final LAX (fig. 5.1f). The estimate for the LAX is used for finding the MVP, in a spherical projection of the LV (fig. 5.1g).

*algorithm
outline*

Detection of long axis | 5.2.1

For the detection of the LAX in the apical 3DE image sequences of the left ventricle we choose to detect the LV centers in planes perpendicular to the acquisition axis. These planes resemble regular 2D short-axis acquisitions of the LV (fig. 5.2a). They are extracted by dividing the image data into a number of slices L and integrating the data within each slice along the acquisition axis. We exclude the upper 10% and the lower 10% of the image volume from analysis to remove influence of near-field noise and to exclude the lower part of the image that has been left empty, respectively.

We aim at detecting the center of the endocardial border, which appears in these planes as the inner edge of the main circular structure, using a Hough transform [Ballard 1981] for circles (HT_c). The Hough transform is known to be robust to partial dropout of target structures and invariant to the circle radius.

*Hough
transform for
circles*




Figure 5.1: The detection scheme for the long axis (LAX) and the mitral valve plane (MVP). *a*) From the original 3D + T image (1), projection slices (2) are extracted at a certain number of levels per cardiac phase. *b*) A Hough transform for circles computes a circle center probability map for each slice (the accumulator image, 3). *c*) For each cardiac phase, 3D dynamic programming determines the LAX path (4) through the probability maps. *d*) The probability maps are weighted according to the detected circle center from the previous step. *e*) The circle centers are tracked through the weighted probability maps (5) over time, per slice level. This results in the circle center paths (6). *f*) For each cardiac phase, a weighted line fit determines the LAX (7). *g*) The LAXs are used to detect the MVP in each cardiac phase (8).

For location of possible circle centers, the HT_c utilizes the image gradient G as a measure for edge orientation and the gradient magnitude $\|G\|$ as a measure for edge strength. G is implemented as a convolution with a Gaussian derivative at a scale σ .

We apply a thresholding operation on the gradient image G to remove influence from noise in the background and use only the strong gradients. We define the threshold as the gradient value belonging to a certain percentile value g_t of the combined histogram of the total set of gradient images to be invariant to global contrast changes throughout the different acquisitions. The optimal value of g_t is determined experimentally. The HT_c transforms the gradient image into a probability map for circle centers, the accumulator image A (fig. 5.2b), with the same dimensions as G , using the gradient magnitude $\|G\|$ as a weighting function for the edge responses, which reduces the sensitivity to the threshold value g_t ,

gradient image

$$A_p = \sum_q (g_{p,q} r_{p,q} \|G_q\|) \text{ where} \quad (5.1)$$

$$p, q \in G$$

$$g_{p,q} = \begin{cases} 1 & \angle(G_q(\vec{p}q)) \leq \alpha_\epsilon \\ 0 & \text{otherwise} \end{cases}$$

for any positions p and q in the gradient image. The parameters r_{min} and r_{max} define the minimum and maximum radius that should be possibly detected by the Hough transform. We set $[r_{min}, r_{max}]$ to $[10\text{mm}, 30\text{mm}]$. Note that the maximum radius r_{max} corresponds well with normal values for the end diastolic (ED) LV diameter (95% interval: 37 - 56 mm, Feigenbaum [2004]). The angle uncertainty α_ϵ is related to the precision of the gradient estimation. It determines the width of the accumulator region for which values are increased. Radius image R (also with the same dimensions as G) accumulates the candidate radii that are detected for a certain position p ,

probable radius

$$R_p = \sum_q (g_{p,q} r_{p,q} \|G_q\| \|\vec{p}q\|). \quad (5.2)$$

An estimate for the most likely radius $\hat{r}_p$ of a circle at p is defined as $\hat{r}_p = R_p / A_p$. We employ this circle detector in L planes perpendicular to the (apical) acquisition axis in all cardiac phases and find probabilities for p being a circle center. These planes in all cardiac phases constitute a 3D plus time (3D+T) probability map for circle centers.




Figure 5.2: LAX detection in a 3D image using the Hough transform for circles (HT_c). The top row shows the accumulator images. In the bottom row the accumulator images are blend in with the original image. *a)* Original image slice perpendicular to image axis. *b)* The HT_c assigns circle center probabilities to the original image slice in the accumulator image. *c)* Center line detection is performed using MDP to find a path approximation of the LAX in 3D. *d)* The accumulator image is weighed using the detected center in each slice. *e)* The circle center trace is detected over time using MDP. *f)* The LV center as detected in the previous step. This LV center will be used for the line fit in each phase.

5.2.1.1 | Dynamic programming

Given the 3D + time probability map, we detect the LAX over the full cycle by multi-dimensional dynamic programming. Dynamic programming (DP) is a well-known graph search technique [Bellmann 1965]. In image processing it is often referred to as a minimum cost algorithm for finding a connective path through a 2D cost image [Amini et al. 1990; Sonka et al. 1999]. In this classical approach (fig. 5.3), the pixels in the cost image v (of M rows and N columns) act as nodes n_{ij} ($i = 0 \dots M - 1$; $j = 0 \dots N - 1$) in a directed graph, the corresponding pixel values v_{ij} as the node costs c_{ij} (i.e. the circle center probabilities). The directional edges in the graph are defined by imposing a connectivity constraint, the maximum step size S . This step size limits the number of neighbors $(2S + 1)$ to which a node n_{ij} in column j is connected in the next column $j + 1$. We denote the directed edge e_{ijk} , the edge from n_{ij} to $n_{i+k, j+1}$, where $k \in [-S, S]$. Additional costs a_{ijk} may be assigned to edge e_{ijk} .

Dynamic programming is a technique that greedily searches to find the cheapest path from column 0 to $N - 1$. This path is found by computing the cumulative

graph search

step size

cumulative costs




Figure 5.3: Dynamic programming. *a)* Initial cost matrix. The costs c_{ij} for each node n_{ij} are shown in the top-left of each node. *b)* Computation of the cumulative cost matrix. Cumulative costs w_{ij} are shown in bold for the nodes in the first two columns. The arrows represent the edges e_{ijk} between the nodes. The bold edge is the edge for which the cumulative costs for the bold node are minimal. *c)* The complete cumulative cost matrix with for each node a reference (representing k) to the previous column. The optimal path is backtracked from the node with lowest cumulative costs in the last column, following the references (redarrows).

cost w_{ij} for each node n_{ij} ,

$$w_{ij} = \min_{k=-S \dots S} (w_{i-k, j-1} + \alpha_{ijk} + c_{ij}) \text{ where} \quad (5.3)$$

$$w_{i0} = c_{i0}$$

The value of k for which w_{ij} is minimal is stored with each corresponding node. *back tracking*
The optimal path is then easily found by backtracking from the node $n_{i, N-1}^*$, for which

$$w_{i, N-1}^* = \min_{i=0 \dots M-1} (w_{i, N-1}) \quad (5.4)$$

by following the edges for the corresponding k -value stored with each node, down to $n_{i,0}$.

This algorithm can be extended to find an optimal path through a multidimensional image (with dimension $D > 2$), by allowing side steps in $D - 1$ dimensions. This extension is known as multidimensional dynamic programming (MDP), as previously presented by Üzümcü et al. [2006].

5.2.1.2 | Continuous LAX detection

*spatial
continuity*

*temporal
continuity*

We aim at finding the LAX in 3D, continuously over time, in the 3D + T accumulator image (fig. 5.2). Therefore, we detect LV centers using two different MDP steps. First we approximate the LAX in each cardiac phase by finding a path in the 3D accumulator image using MDP (i.e. 2D DP in this case). The MDP will obtain a continuous path as an approximation of the LAX for each separate phase (fig. 5.2b,c), but temporal continuity is not imposed in this way. Accordingly, we combine the results from the single-phase 3D MDP detections, with 2D+T MDP detections (i.e. detecting the trace of the LV center at a certain (short-axis) level over time), by weighting the accumulator image with a distance function (fig. 5.2d). The value of this Gaussian distance function decays with the distance from the detected path in 3D. We use the weighted accumulator image as the cost image for the 2D+T detection. In this way, we exploit the continuity along the acquisition axis from the 3D detection and find a continuous LV center path over time for each level (fig. 5.2e). Finally, we employ a weighted least squared distance line fit on the detected LV centers (fig. 5.2f) of the 2D+T MDP, with the accumulator value as the weight, for location of candidate circle centers. The maximum side step S , an integer value, should not be chosen too small to allow enough curvature in the detected path. On the other hand, choosing S too large degrades the computational efficacy of the method and weakens the continuity of the detected path. From our experiments we found that $S = 2$ gives both for the 3D and for the 2D+T MDP enough freedom to find the desired LV center points, without degrading the continuity.

5.2.2 | Mitral valve plane detection

*image
projection*

For the detection of the mitral valve plane (MVP) we use the LAX position from the previous step. Given the LAX estimation in each cardiac phase we estimate the MVP by detecting it in a spherical integration of the LV on a plane through the LAX (fig. 5.4a). We define the MVP as being the plane perpendicular to the detected LAX, touching the bottom of the LV endocardial border. In this spherical projection we obtain a simplified and integrated representation of the data, assuming that the LAX lies within the mitral valve ring and points to the apex, and assuming an approximately ellipsoidal shaped LV. We define a spherical coordinate system (ρ, ϕ, θ) , with ρ for radius, ϕ for elevation and θ for azimuth. The coordinate system has the LAX as its vertical axis. The origin is defined as the weighted center of gravity of the LV centers from the LAX detection. We employ a mean projection of the intensities $I(\rho, \phi, \cdot)$ on a plane $l : \theta = c$. The intensities of the resulting projection image P_c are

defined as

$$P_c(\rho, \phi) = \frac{1}{2\pi} \int_0^{2\pi} I(\rho, \phi, \theta) d\theta. \quad (5.5)$$

This projection exploits the circular shape of the myocardium in short-axis planes, and therefore increases blood-to-tissue ratio (fig. 5.4a,b). We aim at detecting the approximate, projected endocardial border and the MVP in this projection using a low-level edge detection technique. We employ dynamic programming using the radial gradient in the spherical projection as the cost function. A typical result for the projection and the detected border are shown in fig. 5.4c. The detected border is then back transformed into the Cartesian image domain to extract the MVP (fig. 5.4d).




Figure 5.4: *a)* Spherical projection of the image intensities onto a plane. *b)* The projection image, mirrored in the LAX. *c)* The radial gradient of the projected image, with the detected border. *d)* An illustration of the detection of the MVP, using the detected path (in a cylindrical projection (r, d)).

Image acquisition | 5.2.3

Transthoracic apical real-time 3DE images were acquired using the Fast Rotating Ultrasound (FRU) transducer [Voormolen et al. 2006], connected to a Vingmed Vivid FiVe (GE Vingmed, Horten, Norway) and using the commercially available Philips

*FRU
transducer
data*

Sonos 7500, with the X4 xMatrix transducer (Philips Medical Systems, Andover, Massachusetts, USA). The FRU acquisitions were made on a group of 11 patients (age: 52 ± 12 years), with a diagnosis of myocardial infarction (MI). These patients were selected from an initial group of 14 patients, which were included based on sufficient 2D echo image quality. Three patients had severely dilated ventricles that could not be imaged entirely and were therefore excluded from the study. Acquisitions were made 156 ± 82 days after MI. Image sequences were interpolated to volumetric data (chapter 4), $256 \times 256 \times 400$ pixels at 16 phases per cardiac cycle. For automated analysis the sets were downsampled to $128 \times 128 \times 400$ pixels, to reduce computational costs without degrading the results.

*matrix
transducer
data*

Another group of 14 patients (age: 57 ± 14 years), who were referred for dobutamine stress echo, were examined using the Philips Sonos 7500. These (rest) data sets varied from 15 to 24 phases per cardiac cycle and had dimensions of $144 \times 160 \times 208$ pixels. Both acquisition sets contained image sequences of varying image quality. An example of both types of patient data (of average image quality) is shown in fig. 5.5.

5.2.4 | Evaluation

Both the FRU and matrix acquisitions were analyzed manually using a semi-automatic segmentation tool for quantitative assessment of full cycle LV volumes (chapter 2). Two observers analyzed FRU data independently, after reaching agreement on the tracing conventions. The matrix acquisitions were traced independently by another observer.

*semi-
automatic
endocardial
contour tracing*

Full cycle endocardial contours were traced semi-automatically by drawing contours in four 2D intersections per patient, followed by automatic detection. If desirable, corrections were made iteratively to achieve a fully satisfying segmentation in all the cardiac phases. The endocardial contours from these tracings were used to determine the manually defined LAX and MVP. We derived two different long axes from the manual segmentations. Generally the LAX is defined as the line segment between the mitral valve center (MVC) and the point on the contour with largest distance to the MVC. We define our regular LAX (rLAX, fig. 5.6) as the line segment from the MVC to the center of gravity (COG) of the apical volume (top 25%), to be less sensitive to small irregularities in the apical contour. A disadvantage of these definitions is that it may result in a rLAX that is intuitively off-center for a bent LV. Therefore we also compute a centerline LAX (cLAX, fig. 5.6), which is a line fit through the short-axis (given the rLAX) contour centers. Note that for automatic initialization purposes the actual definition of the LAX is not critical as long as it represents the main shape, is robust to small contour changes and can be esti-

regular LAX

centerline LAX




Figure 5.5: Two examples of the patient data, note the deviation of the acquisition axis from the LV long axis. *top*) Three orthogonal slices of a FRU data set of average image quality. *bottom*) Three orthogonal slices of a Philips Sonos 7500 (matrix transducer) data set of average image quality.

mated accurately. We define the manual MVP as the least squares plane fit through the mitral valve ring points. Note that it is not necessarily perpendicular to the LAX.

We evaluate the distance and angle of the detected LAX to the rLAX and the cLAX. The distance is defined as the smallest Euclidean distance between the manual LAX line segment and the automatically detected LAX line segment in mm. Such a distance by itself is not a very discriminative measure for evaluation of the quality of detected axes. Two axes may be almost intersecting, and thus have a small distance from each other, but may point in a totally different direction. Therefore, we also measure the angle between the vectors belonging to these axes (in degrees). If both the distance and the angle are small, the axes are similar.

We measure the quality of the detected mitral valve plane as the projected signed distance between the detected MVC and the manual MVP. A negative distance means that the detected MVP lies above the manual MVP (within the LV cavity). We measure the distance, because for initialization purposes we are mostly interested in the level of the MVP, not in the angle with respect to the LAX. The proposed method

*LAX error
measure*

*MVP error
measure*




Figure 5.6: The definition of the regular long axis (rLAX) and the centerline long axis (cLAX).

does not measure this angle, because it assumes that the MVP is approximately perpendicular to the LAX.

We optimized the performance of the LAX and MVP detection by systematically varying the free parameters of the method for both types of acquisition systems individually, for the complete cardiac cycle. The metric for this optimization is composed of the measured mean and standard deviation for distances and angles of the detected LAX to the manual LAX (rLAX or cLAX). In this metric, distances and angles are normalized according to the found interobserver variabilities (see below).

parameter optimization For initial estimation of the optimal parameters we assumed them to be independent. We evaluated the following parameter ranges and increments ($\{\text{parameter}; \text{range}; \text{increment}\}$): $\{\sigma; [0.5, 3.0]; 0.5 \text{ sd}\}$, $\{g_t; [70, 95]; 5\%\}$, $\{L; [5, 30]; 5\}$ and $\{a_c; [5, 40]; 5^\circ\}$ for both acquisition methods. After determining probable ranges, we optimized the parameters, without assuming independence, thus by full exploration of the determined remaining parameter space.

5.3 | Results

5.3.1 | Interobserver variability

We determined interobserver variabilities from the manual FRU tracings by two observers. These interobserver variabilities were determined for the full cardiac cycle (table 5.1 and fig. 5.7). We also obtained interobserver variabilities for the

distance between the observer's MVPs over the full cardiac cycle (fig. 5.9). The average (signed) interobserver distance was 0.94 ± 1.80 mm, with a point-to-point distance for the MV center of 3.65 ± 1.83 mm. Note that interobserver variabilities found here are lower than can be expected from a range of users from different institutions, because both observers reached consensus on the tracing conventions, before analyzing the patient data.

Table 5.1: Interobserver variabilities for two observers on FRU data of 11 patients ($N = 176$ frames).

	Distance (mm)	Angle ($^{\circ}$)
rLAX	1.39 ± 1.07	3.40 ± 1.72
cLAX	1.29 ± 0.98	3.15 ± 1.78

All the results are expressed as mean $\pm$ standard deviation.

Parameter optimization | 5.3.2

After initial (independent) parameter optimization for all the parameters $\{\sigma, g_t, L, \alpha_e\}$ for each of the acquisition methods, we determined smaller ranges and smaller step sizes for the full (dependent) optimization. For the FRU data we found the following ranges and step sizes: $\{\sigma; [0.5, 1.5]; 0.5\text{sd}\}$, $\{g_t; [90, 97.5]; 2.5\%\}$, $\{L; [11, 19]; 2\}$ and $\{\alpha_e; [10, 30]; 5^{\circ}\}$. For the matrix data these were: $\{\sigma; [0.5, 2.0]; 0.5\text{sd}\}$, $\{g_t; [85, 95]; 2.5\%\}$, $\{L; [11, 19]; 2\}$ and $\{\alpha_e; [10, 30]; 5^{\circ}\}$. The full exploration of these acquisition dependent parameter spaces resulted in the optimal parameters for the LAX detection (table 5.2). The detection results after full optimization of the parameters improved only a few percent with respect to initial (independent) parameter optimizations. This shows the relatively low sensitivity of the detection to small parameter changes.

We found very similar parameters for the two data types. This can be attributed to the fact that the algorithm detects a very coarse structure. At such a scale, differences between both acquisition systems are small. The optimal parameters differ most for σ . This may be due to the lower azimuth resolution of the FRU transducer, which results in a higher σ .

Table 5.2: The optimal parameter settings for LAX detection in FRU and matrix data.

	rLAX				cLAX			
	σ (pixels)	g_t (%)	L (#)	α_e ($^\circ$)	σ (pixels)	g_t (%)	L (#)	α_e ($^\circ$)
FRU	1.0	90	15	25	1.0	90	15	25
Matrix	0.5	92.5	15	35	0.5	92.5	15	30

5.3.3 | LAX detection

Initial LAX errors (with the acquisition axis as LAX estimate) and detection results for FRU and matrix data are shown in table 5.3 and 5.4 respectively. LAX detection results improve significantly for all cardiac phase (distances and angles) with respect to the initial errors ($p < 0.01, N = 25$), for both FRU and matrix data (fig. 5.8). But also significant differences are found between interobserver variabilities and detection errors in some cardiac phases ($p < 0.05, N = 11$). Nevertheless, detection errors are small and acceptable for initialization purposes, if compared to expected clinical interobserver variabilities.

Detection results are comparable for both acquisition types. FRU data yields slightly lower angle errors, while matrix data yields lower distances. Overall, the LAX detection approximates the cLAX better than the rLAX, although differences are small. This is to be expected, as our LAX detection scheme resembles the computation of the cLAX.

Table 5.3: Initial and detection results for LAX on FRU data ($N = 176$ frames). The distance and angles are computed with respect to the manual rLAX and cLAX.

	rLAX		cLAX	
	Distance (mm)	Angle ($^\circ$)	Distance (mm)	Angle ($^\circ$)
Initial	5.41 $\pm$ 3.54	11.23 $\pm$ 5.30	6.38 $\pm$ 3.36	10.75 $\pm$ 5.37
Detected	2.85 $\pm$ 1.70	5.25 $\pm$ 3.17	2.32 $\pm$ 1.49	4.76 $\pm$ 2.95

All the results are expressed as mean $\pm$ standard deviation.




Figure 5.7: Interobserver variabilities for FRU data for the LAX annotation. The mean and standard deviation for the distances are plotted for the rLAX and cLAX, with their corresponding maximum values (dashed line).




Figure 5.8: Detection results for the LAX (FRU and matrix combined). The mean and standard deviations for the angles are plotted for the rLAX and cLAX, with their corresponding maximum values (dashed line). (†) denotes a significant difference ($p < 0.05$, $N = 25$) between initial and detected errors for the cLAX. (‡) denotes a significant difference ($p < 0.05$, $N = 11$) between the interobserver variability and the detected error (FRU data only, cLAX).

Table 5.4: Initial and detection results for LAX on matrix data ($N = 224$ frames). The distance and angles are computed with respect to the manual rLAX and cLAX.

	rLAX		cLAX	
	Distance (mm)	Angle ($^{\circ}$)	Distance (mm)	Angle ($^{\circ}$)
Initial	3.12 ± 2.15	11.17 ± 3.48	3.90 ± 2.45	10.38 ± 3.46
Detected	1.96 ± 1.30	5.95 ± 2.11	1.87 ± 1.28	4.96 ± 1.93

All the results are expressed as mean $\pm$ standard deviation.

MVP detection | 5.3.4

MVP detection results for FRU and matrix data are listed in table 5.5 and fig. 5.9. Small mean errors with low standard deviations were found for both data types, although the automated detection significantly underestimates the depth of the MVP in most systolic phases ($p < 0.05$, $N = 11$). In comparison to interobserver variabilities, errors are especially higher in systole, where the automated detection underestimates the displacement. Near end diastole (ED, phase 1) and end systole (ES, $\approx$ phase 8) errors are smallest and differences with the interobserver variabilities are not significant ($p > 0.05$, $N = 11$).

Point-to-point distances are significantly higher for the automated method compared to interobserver variability in most phases ($p < 0.05$). In this measure, the error in the LAX estimation is reflected, because the detected MV center is based on the estimated LAX.

Table 5.5: MVP detection results for FRU and matrix data.

MVP	FRU (176 frames)	Matrix (224 frames)
Point-to-Plane (mm, signed)	3.12 ± 2.15	11.17 ± 3.48
Point-to-Point (mm)	1.96 ± 1.30	5.95 ± 2.11

All the results are expressed as mean $\pm$ standard deviation.




Figure 5.9: MVP interobserver variability (FRU) and detection results (FRU and matrix combined). *top*) Mean point-to-plane distances (MV center to MVP) and SD over the cardiac cycle. A positive error for the detection results means an overestimation of the MVP depth. *bottom*). Mean point-to-point distances (MV centers) and SD over the cardiac cycle. (*) denotes a significant difference ($p < 0.05$, $N = 11$) between the interobserver variability and the detected error (FRU data only).

Discussion | 5.4

We presented a method for automatic detection of the LV LAX and the MVP over the full cardiac cycle. It is based on a Hough transform for circles for finding LV center probabilities in integrated slices perpendicular to the (apical) acquisition axis. Subsequently multi-dimensional dynamic programming is used to detect circle centers using these probabilities, continuously along the LAX and over time. We employ least squared distance line fit to find the LAX in each cardiac phase. The MVP is located using the detected LAX in a spherical projection, using low-level edge detection by dynamic programming.

The method consists of a few consecutive steps. After determination of LV center probabilities using HT_c on integrated slices of the volume, MDP in 3D results in continuous LV center probabilities and serves as outlier removal for the LV center probability. The same holds for MDP over time. In this way, all available data is combined and a robust continuous detection is achieved.

With respect to the initial LAX errors, the method shows significant improvement, with maximum distance errors of 6.07 mm in ED and 5.07 mm in ES, and angle errors of 8.80 for ED and 8.75 for ES (all for the cLAX). These maximum errors indicate the usability of the LAX detection as an initialization step, especially, when taken into account that ED and ES are the most important cardiac phases for initialization of automated procedures, because they represent the two geometrically extreme states of the LV. Moreover, LAX detection did not show significant differences with respect to our interobserver variabilities, which were achieved in the idealized situation where both observers agreed on endocardial border tracing conventions in advance. In daily clinical situations, these observer variabilities are expected to be considerably larger.

*maximum
errors*

In the analysis of the interobserver variabilities and the detection errors we normalized the patients' cardiac cycles to 16-phases cycle. In this normalization the position of the ES-phase has not been taken into account due to small variation in the duration of systole among the patients. Therefore, in the analysis shown in 5.8 and fig. 5.9 there is no cardiac phase that can be depicted as ES. Instead, ES has been faded over a few cardiac phases.

*definition of
end systole*

The detection of the MVP depends on the detection of LAX. This may influence automatic location of the MVP in patients suffering from pathologies that alter the LV shape considerably. The method may be extended by applying same kind of quality control or reliability estimate e.g. by using the quality of the final line fit in estimating the LAX per cardiac phase. Nevertheless, small deviation of the LAX from its true position, does not affect the MVP detection much. This is because the mitral valve is approximately a planar structure, almost perpendicular to the LAX.

*LAX depends
on MVP
detection*

The detection of the MVP is limited to finding a MVP plane perpendicular to the LAX, while the true MVP usually is not located exactly perpendicular to the LAX. We found deviations of 5.18 ± 2.81 degrees (mean $\pm$ standard deviation) from the plane perpendicular to the LAX ($N = 400$ frames) in our manual tracings. These small differences may be discarded for our initialization purposes.

Automatic initialization of automated segmentation of the LV in 3DE decreases analysis time for assessment of LV function drastically. Moreover, it eliminates observer variability and therefore makes measurements more reproducible and therefore allows inter-institutional comparison of LV function assessments. This is of great value for large studies, for example in clinical trials. However, it of course remains to be proven that such automated measurements, employing a combination of automated analysis and detection, is accurate in comparison to the gold standard.

*exploration of
left ventricle*

The presented method provides a basis for localization of LV salient structures, as has been illustrated by the detection of the MVP. Given the location and orientation of these landmarks one can estimate the complete orientation of the LV using the knowledge about the acquisition to determine the angle of the right ventricle (RV) with respect to the LAX. This makes it a suitable method for initialization of subsequent processing steps, such as LV segmentation.

apex detection

As a complete initialization approach for LV segmentation, the proposed method lacks true apex detection. This should not be seen as an important shortcoming of the method. We expect that initialization can be done reliably using the LAX and MVP. The remaining freedom in the LV position and orientation is very limited, and final determination of this position, orientation and shape should be treated by the segmentation approach.

5.4.1 | Study setup

pathologies

In our experiments we evaluated the LAX and MVP detection on data from patients with various diagnoses of cardiovascular disease. In this population the method showed robustness combined with good accuracy. Although patient data was used for evaluation, the method might encounter problems in very pathological cases showing aberrant LV shapes (e.g. apical aneurysms). These topics need to be further investigated. Note however, that the definition of the LV LAX in these cases also is problematic. Largest errors in the LAX detection in our study (which accounted for the maximum angle errors in systole, fig. 5.8) were caused by one patient with a highly trabeculated ventricle, which fooled the LV center detection. This would be a subject for further research on LAX detection.

The application to patient subpopulations, such as patients with an extremely

dilated LV probably requires adjustments of the method's parameters concerning the expected LV diameter ($r_{\max}$).

In this study, because of the limited availability of patient data, the same data sets were used for parameter training (optimization) as for testing, while ideally these data sets should be different. However, the parameter optimizations have shown that parameter choices are not very critical, because of the very small gain in performance when parameters were optimized with dependence among them. In this optimization, where the optimum neighboring parameter space has been fully explored, in 80% of the evaluations the objective metric was less than 12% above the optimum. This range is small compared to the initial estimation where the objective metric is 96% higher. Also, differences in optimal parameters between the two acquisition methods were small. In an evaluation of the LAX detection on FRU and matrix data with the mean of the individually optimized parameters (fig. 5.2), the objective metric deviated less than 1% from the optimal case. For these reasons, very similar results can be expected if the training set is separated from the test set.

*bias in
optimization*

Performance of matrix vs. FRU data | 5.4.2

The initial distance errors for the LAX are lower for the matrix acquisitions. This can be attributed to the difference in the acquisition procedure. For the matrix system, a bi-plane view is shown when positioning the probe. For the FRU acquisitions, the sonographer currently needs to alter between the two planes, which makes positioning slightly more difficult. This is reflected in the higher initial distance errors for the FRU data.

*difference in
acquisition*

Automatic detection yields very similar performance for the two acquisition types. The differences in distance and angle errors between them were not significant ($p > 0.05$). Detections on matrix data show lower distance errors, while detections on FRU data show a slightly lower angle error. These differences are small and might be due to differences in the patient populations.

*no significant
difference*

Hough transform for circles | 5.4.3

The analysis of the LV in planes perpendicular to the acquisition axis assumes a circular shape of the myocardium in these planes. This is an important assumption in the estimation of the LV centers using the Hough transform for circles. However, either due to deviation from the true short-axis angle (in fact we are trying to detect this angle) or due to the left ventricular shape, the LV may appear like an ellipse in

*extending the
Hough
transform*

the images in which the LV center is detected. As a consequence, one might argue that a Hough transform for ellipses would be more appropriate when detecting the LV center. A disadvantage of a Hough transform for ellipses would be that two extra parameters must be estimated, namely the minor radius and the rotation angle of the ellipse. This would make the detection computationally more expensive. The same holds for the possible extension of the Hough transform to 3D for detecting ellipsoids.

In practice, the HT_c will also be capable of approximating centers of ellipses with arbitrary minor radius and orientation, as long as the major and minor radius are close to, or within the accepted range of the radius for the HT_c . Also, the actual deviation of the planes perpendicular to the acquisition axis from the true short-axis planes is limited as the initial errors show (table 5.3 and 5.4), because in the acquisition the sonographer aims at aligning the acquisition axis to the LV LAX.

5.4.4 | LV apex detection and LAX length

In the detection of the MVP, the projected endocardial border is detected using the radial gradient in a spherical projection. This border is used to find the MVP. Similarly, the apex could be detected for computation of the LAX length. The highest point of the detected border, or the intersection of this border with the LAX could be used as estimation for the LV apex. A drawback of such an apex detection is that it is very sensitive to the initially detected LAX. A small deviation of the LAX from its true position leads to a considerably lower intersection of the LAX with the endocardial border, especially when the ventricle is narrow near the apex, resulting in a misplaced apex and an underestimation of the LAX length. Furthermore, the presence of near field artifacts in the image obscures the apical region in the projection image. This complicates the detection of the apex in these images using a low-level border detection technique. Therefore, we leave the task of detecting the apex to the proper segmentation method.

5.4.5 | Extensions

right ventricle

The presented method detects the LV LAX, and using this LAX, also the MVP. Once the LAX is defined, a multitude of possibilities for detecting other salient structures may become feasible. One of the most desirable structures may be the right ventricle (RV), or more specifically, the RV attachment point. This would allow full determination of the position of the LV in 3D, with respect to all six degrees of freedom (translation and rotation). The presented method leaves the rotation around

the LAX (z -rotation) unattended. Note however, that the variation in z -rotation is limited, because acquisitions are made using the 2- and 4-chamber view as a reference. The detection of the RV in data from current 3DE scanners is problematic, because in many regular acquisitions the RV is hardly visible. The (anterior) RV attachment point is visible in most cases, but the occasional absence of this point and the presence of image artifacts that may fool the detection, making it a feature that is hard to locate fully automatically. When feature detection is used for automatic initialization purposes, such failures are very undesirable.

Currently, we detect the level of the MVP with respect to the LAX. This serves its goal as an initialization for segmentation of the LV. A desirable extension of the method may be automatic detection and tracking of the MV hinge points (MVHP). The mitral annular motion is useful in the evaluation of global and regional LV function and an important parameter in the diagnosis of annular diseases and LV disorders [Eto et al. 2005; Pai et al. 1991; Willenheimer et al. 1999]. The MVHP are typically visible as bright structures and seem suitable for automated detection. Automated tracking of the MVHP in 2D echocardiography has been shown to be feasible in chapter 3.

Computational costs | 5.4.6

For initialization methods, low computational costs are obviously desirable. The detection of the LV LAX and MVP over the full cardiac cycle (16 phases) took two to four minutes on a regular PC (Intel Pentium IV, 2.6 GHz), depending on parameter choices. The implementation of the method (in C++) was not optimized for speed and is suitable for parallel processing. Furthermore, the method may be considerably sped up by applying it at a lower resolution because of the coarse nature of the desired feature detection. Besides, the optimization of the method's parameters shows that parameter choices are not very critical. This gives room for parameter choices that increase performance in terms of computational costs, without noticeably decreasing the accuracy of the method.

Conclusions | 5.5

We presented a method for automatic detection of the LV LAX and the MVP over the full cardiac cycle. It is based on a Hough transform for circles and multidimensional dynamic programming for detecting the LV LAX continuously over the cardiac cy-

cle in 3DE. Using the detected LAX, it locates the MVP by employing a DP border detection in a spherical projection of the 3D LV. In an evaluation on FRU and matrix data, the method has shown to be robust and accurate in detecting the LAX and MVP. The accuracy, combined with its low computational costs, make it very suitable for initialization purposes for automated segmentation algorithms for the LV.

| Acknowledgments

We gratefully acknowledge B.J. Krenning, F.J. ten Cate and M.L. Geleijnse (Thoraxcenter, Erasmus MC Rotterdam) for providing the patient data that was used in the study.

Automated left ventricular volume estimation in 3D echocardiography using active appearance models

ASSESSMENT OF LEFT VENTRICULAR (LV) functional parameters, such as LV volume, ejection fraction and stroke volume, from real-time 3D echocardiography (3DE) is labor intensive and subjective, because in current analyses it requires input from the user. Automating these procedures will save valuable time in the analysis and will remove interobserver variability.

We investigated a fully automatic segmentation approach for the left ventricle in real-time 3D echocardiography, based on active appearance models (AAMs). AAMs were built with end-diastolic images from 54 patients. We evaluated generalization capabilities of the shape and texture model and matching performance of the AAM using regular and Jacobian tuning matching algorithms in various scenarios.

The generalization of the shape model was good, comparable to a model containing 97% of the total modeled variation. The generalization of the texture model was moderate, comparable to a model containing 70% of the variation, which may hamper the AAM matching. In the comparison of the regular and Jacobian tuning matching methods, the latter obtained larger capture ranges and a higher accuracy.

The matching results indicate that fully automatic segmentation of the LV in 3DE using AAMs is feasible. Jacobian tuning matching has shown great potential for segmentation in echocardiograms and will improve the assessment of LV functional parameters.

This chapter is partially based on:

Automatic segmentation of the left ventricle in 3D echocardiography using active appearance models. M. van Stralen, K.Y.E. Leung, M.M. Voormolen, N. de Jong, A.E.W. van der Steen, J.H.C. Reiber, J.G. Bosch. Proc IEEE Int Ultrason Symp 2007; 1480-1483 (© 2007 IEEE) and **Improving 3D active appearance model segmentation of the left ventricle with Jacobian tuning.** K.Y.E. Leung, M. van Stralen, M.M. Voormolen, N. de Jong, A.E.W. van der Steen, J.H.C. Reiber, J.G. Bosch. Proc SPIE Med Imaging 2008; 6914; 69143B (© 2008 SPIE)

6.1 | Introduction

6.1.1 | Goal

Assessment of left ventricular (LV) functional parameters, such as LV volume, ejection fraction and stroke volume, from real-time 3D echocardiography (3DE) is labor intensive and subjective, because in current analyses it requires input from the user. Automating these procedures will save valuable time in the analysis and will remove interobserver variability.

Previous chapters of this thesis have focused on semi-automated approaches for automated 3D segmentation of the left ventricle, and preprocessing steps. The current chapter is devoted to our achievements towards the realization of a fully automated 3D segmentation approach based on so-called active appearance models (AAMs). AAMs hold considerable promise for the difficult task of segmentation in 3D ultrasound, and many research groups have addressed the topic, but the issue has not been solved so far.

6.1.2 | Related work

Previously, various techniques for the automated analysis of the left ventricle have been presented [Angelini et al. 2005; Corsi et al. 2002; Gérard et al. 2002; Hansegård et al. 2007a; Kühl et al. 2004; Zagrodsky et al. 2005; Zhu et al. 2007], nevertheless, most of these [Angelini et al. 2005; Corsi et al. 2002; Gérard et al. 2002; Kühl et al. 2004; Zagrodsky et al. 2005] still require manual interaction in the form of some indicated landmarks or manually drawn contours to achieve a proper analysis. Zagrodsky et al. [2005] have presented a fully automatic segmentation approach for the LV in 3DE, which is initialized by registration of a presegmented template with the unseen image. This approach is time-consuming and requires a presegmented template that can be successfully matched to any image to initialize the segmentation method. Furthermore, the evaluation was done on a limited number of subjects and showed significant problems when the ventricle was not fully captured in the imaging volume. Hansegård, Orderud et al. [Hansegård et al. 2007b; Orderud et al. 2007] combine an active shape model with a Kalman filter and show promising results. Zhu et al. [2007] attempt to detect the endocardial and the very challenging epicardial border using a maximum-a-posteriori framework which incorporates a statistical speckle model and an incompressibility constraint for the myocardium. Nillesen et al. apply adaptive filtering using image statistics as a preprocessing step to the automated segmentation [Nillesen et al. 2007; Nillesen et al.

2008]. The method lacks a constrictive shape model which hampers the detection.

Extensive research on analysis of time varying 2D echocardiographic images has been carried out by Comaniciu et al. [2004] and Jacob et al. [2002]. They integrate temporal, textural information with an adaptive shape model in a Kalman filter approach (see also section 6.5.5.1).

Fully automated 3D LV segmentation approaches in other modalities like CT [Zheng et al. 2007] and MR [Kaus et al. 2004] do not translate well to ultrasound. Mostly, these rely strongly on the generic intensity difference of tissue and blood and only apply weak shape continuity constraints on the deformable model (balloons, level sets etc). Due to the significant artifacts in ultrasound (in particular drop outs and clutter), anisotropy and position and orientation dependent intensity characteristics, ultrasound approaches require stronger shape constraints and more localized modeling of appearance. *other modalities*

Proposed approach: active appearance models | 6.1.3

Most segmentation and analysis approaches take a data-driven or bottom-up approach: they derive features from image data and try to fit a model (geometric, patterns, etc.) to these features. In our case, we propose to follow an opposite approach, labeled analysis-by-synthesis.

We want to analyze a complex object that has a well-defined topology, but can exhibit a wide range of natural variability in shape and intensity patterns. Suppose we have a way of synthetically generating realistic images and can cover the variability with a limited number of parameter settings; then we can solve the analysis problem by finding the parameters that generate the best-fitting image. An example of such an approach is that of active appearance models. *analysis-by-synthesis*

We aim at fully automatic segmentation of the left ventricle in 3DE using active appearance models (AAMs). AAMs were first introduced by Cootes and Taylor [2001b], as extension of active shape models (ASMs). This approach has proven to be successful in various image segmentation tasks, starting with face recognition and later on in medical image segmentation. Bosch et al. [2000] have been exploring the application of AAMs to 2D echocardiography (2DE) and introduced the active appearance motion models (AAMMs) for analysis of 2DE time sequences [Bosch et al. 2002]. Since then, AAMs have been adopted to many medical image segmentation tasks and now have taken an important place in medical image analysis research, see section 6.2.2.

6.2 | Active appearance models

6.2.1 | Basic formulation

AAMs represent the shape and the texture of a certain (part of an) organ (in our case, the left ventricle in 3D ultrasound) as a mean appearance with its eigenvariations, by applying principal component analysis (PCA) on example training data annotated by experts. A short, general description of the AAM framework is given below. A complete description can be found in Cootes and Taylor [2001a].

6.2.1.1 | Model generation

We describe the training samples $i \in \{0 \dots N-1\}$ by their shape $s_i = \{x_0, y_0, z_0, \dots, x_{S-1}, y_{S-1}, z_{S-1}\}$ containing S corresponding surface points $\{x, y, z\}$ and their texture $t_i = \{g_0, \dots, g_{T-1}\}$ containing T corresponding image samples g . By applying PCA on these shape and texture vectors we can describe the shape and texture by their mean shape $\bar{s}$ and texture $\bar{t}$, eigenvector matrix Φ_s and Φ_t and the parameter vector b_s and b_t . This requires a definition of (anatomical) point correspondence between the shapes, aligning them to the same pose and size, and calculating the average shape. This shape alignment step is usually solved through Procrustes alignment [Goodall 1991; Gower 1975]. After that, we can perform PCA on the shapes. By warping all shapes to the average shape, we get voxelwise correspondence over the neighborhood of the shapes, and we can calculate an average texture (voxel set) and apply a PCA on texture as well. This gives us a compact description of both shape and texture:

$$s = \bar{s} + \Phi_s b_s \quad (6.1)$$

$$t = \bar{t} + \Phi_t b_t \quad (6.2)$$

We can combine the shape and texture to model possible correlation between typical shape and texture variations by applying a third PCA on the combined parameter vector

$$b = \begin{pmatrix} W_s b_s \\ b_t \end{pmatrix} \quad (6.3)$$

where W_s corrects for the difference in units between shape and texture, to model the appearance,

$$b = \Phi_c c \quad (6.4)$$

The parameter vector c is often extended with parameters describing the pose in 2D or 3D. Depending on the application and the variability that is allowed in the shape model, most commonly uniform scaling, rotation and translation constitute these pose parameters. The appearance and pose parameters μ are combined in the parameter vector $p^T = (c^T | \mu^T)$.

Model matching | 6.2.1.2

One of the main advantages of an active appearance model over other segmentation strategies is its ability to quickly find a good match to the unseen image. An update strategy is used, which only needs a multiplication of the difference image, the difference between the synthesized and the underlying unseen image, with a precomputed parameter update matrix. AAMs are matched iteratively to unseen data by evaluating the difference between the modeled texture t_m and the corresponding texture in the sample image t_s , the residual vector

$$r(p) = t_s - t_m \quad (6.5)$$

and minimizing $E(p) = r^T r$. Minimization of $E(p)$ is achieved iteratively by trying to minimize $E(p + \delta p)$. The first order Taylor expansion of eqn. 6.5 is

$$r(p + \delta p) = r(p) + \frac{\partial r}{\partial p} \delta p \quad (6.6)$$

where $\frac{\partial r}{\partial p}$ is the Jacobian J . By differentiating $E(p + \delta p)$ to p and equating it to zero, we obtain the RMS solution,

$$\delta p = -Ur(p), \text{ where } U = \left(\frac{\partial r}{\partial p}^T \frac{\partial r}{\partial p} \right)^{-1} \frac{\partial r}{\partial p}^T \quad (6.7)$$

For the derivation of eqn. 6.7 we refer to Cootes and Taylor [2006]. An update of p is thus simply generated by a multiplication of the pseudo-inverse of the Jacobian $\frac{\partial r}{\partial p}$, matrix U . The Jacobian J (and thus also U) is assumed to be constant during the matching process and is estimated once and for all matchings in the model training phase.

Model training | 6.2.1.3

In the model training we learn the relation between each of the model parameters, and the difference image. This is done by perturbing each of the model parameters with predefined step sizes and learning the changes in the difference image that appear. This results in a matrix that describes the relation between model

estimating the
Jacobian

parameters and the difference image (the Jacobian J). During the matching we use the inverse relation (update matrix U) to update the model parameters by simply multiplying the difference image with the update matrix, which speeds up the matching tremendously, compared to classical optimization approaches. We estimate the Jacobian by perturbing the model's parameters for every training sample t_l ($l = 1 \dots N$) as follows:

$$\frac{dr_i}{dp_j} = \frac{1}{N} \sum_l \sum_k w(\delta c_{jk}) \left(r_i(p + \delta c_{jk}) - r_i(p) \right) \quad (6.8)$$

where δc_{jk} is the k^{th} perturbation of parameter p_j and $\sum_k w(\delta c_{jk}) = 1$ for all j . The different perturbations of each parameter p_j are weighed using a weighting function $w(\cdot)$, which is usually a suitably normalized Gaussian weighting function or a uniform weighting function, as in our case.

This estimation assumes that the Jacobian is more or less constant near the global optimum. Various studies have shown that this assumption can be made in varying applications. However, ideally one would want to compute the true Jacobian at each position. This would be a computationally very expensive operation. Lately, Cootes and Taylor [2006] proposed a method that iteratively updates the Jacobian during the matching. We will discuss this technique in section 6.3.2.4.

6.2.2 | Evolution of AAM (organs, modalities, dimensions + history)

The classical AAM as defined by Cootes and Taylor has been extended to a range of applications. Original applications by Cootes and Taylor concentrated on segmentation of 2D images of human faces, with some extensions to various medical imaging subjects, such as analysis of vertebral structure in X-ray images [Roberts et al. 2003; Roberts et al. 2007] and MRI of the brain [Cootes and Taylor 2001b].

Interesting work on the analysis of metacarpals in X-ray images was performed by Thodberg [2002], including handling of occlusions. Mitchell, Bosch et al. applied AAMs on cardiac MRI and ultrasound images, first on single 2D cross sections [Bosch et al. 2000; Mitchell et al. 2001a], later extended to time series (Active Appearance Motion Models or AAMM [Bosch et al. 2002; Mitchell et al. 2001b]). For ultrasound, this involved a nonlinear image intensity normalization to overcome the problem of the highly non-Gaussian gray value distribution in ultrasound. Later on, the first 3D implementation of AAM was realized and applied to end-diastolic 3D cardiac MRI datasets and pseudo-3D ultrasound datasets [Mitchell et al. 2002]. The ultrasound datasets were actually time sequences of 2D 4-chamber images stacked into a 3D block, and represented a cylindrical structure with limited 3D freedom. For 3D ultrasound, therefore, this was merely a proof of principle, not

a full realization of a 3D AAM. At the same time, a 2.5D implementation of AAM was realized by Beichel et al. [2002] for segmentation of the diaphragm in CT images. This application was not a full 3D implementation, since it modeled the object as a set of 2D points with the z -coordinate as an attribute, not as a truly 3D shape. Very interesting work has been realized by Stegmann et al.: application in 2D cardiac MRI [Stegmann et al. 2003], extension to a multi-view cluster-aware AAM on cardiac contrast perfusion MRI sequences ([Stegmann et al. 2005] and to a bi-temporal 3D AAM for automated estimation of the ejection fraction in 3D MRI sequences [Stegmann and Pedersen 2005]. Furthermore, 3D segmentation problems have been mostly approached using multi-plane AAM solutions, i.e. using multiple 2D cross sections, either uncoupled [Üzümcü et al. 2005] weakly coupled [Hansegård et al. 2007a] or coupled [Leung et al. 2006b; Oost et al. 2006; Stegmann et al. 2005]. Such approaches have the benefit of reduced computational load and complexity of modeling and matching. However, they may pose unrealistic constraints on shape change or 3D motion. Oost et al. [2006] relaxed these constraints by employing a dynamic programming detection step using the AAM segmentation.

Inspired by the work of Bosch et al., Hansegård et al. [2007a] have shown that active appearance models can be applied with success in triplane echocardiograms. They applied multi-view and multi-frame active appearance models, and compared unconstrained AAMs with AAMs that were constrained by manually placed markers and by dynamic programming (DP). A DP-constrained AAM proved to work best in this setting.

AAMs in echocardiography

In contrast to our approach, Hansegård et al. used a sparse (triplane) AAM, not a full 3D AAM. Furthermore, only a weak coupling in pose was used, to ensure that scale and vertical position did not deviate much over the three views. Hansegård et al. also employed the nonlinear gray value normalization described earlier by our group.

Several modeling and matching methods have been proposed to generate more robust AAM segmentation results. For example, Gross et al. [2006] developed algorithms to apply AAMs to images of faces with occlusions, by combining their inverse compositional approach with a robust error function. An other robust approach for detecting object pose in stereo images consisted of selecting the appropriate multi-view appearance models and subsequent optimization of the robust error function with a modified Gauss-Newton algorithm [Mittrapiyanuruk et al. 2005]. Beichel et al. [2005] proposed a mean-shift-based method to estimate outlier residuals during the matching process. Their approach was applied to different types of medical images containing large artifacts. Recently, Cootes and Taylor [2006] proposed a new Jacobian tuning method, which allows the model's training matrix to adapt itself to new, unseen images during matching. The method is

robustness of AAMs

supposedly more robust, is comparable with respect to speed with the standard matching method, and requires no extra steps in the model-training phase.

6.2.3 | Motivation for our work

AAMs have successfully been applied to a range of segmentation problems, including ultrasound, where they have been shown to offer significant advantages. However, a true 3D AAM for ultrasound has not been demonstrated yet. This could offer significant benefit over 2D or sparse approaches, and it could be extended in several ways, e.g. into a multi-phase 3D or truly 4D approach, a hybrid approach etc. This formed the motivation of our work, and several of these ideas have been investigated, albeit not always with a definite answer.

6.3 | Methods

6.3.1 | Data acquisition

In this work we used clinical data of 54 patients, acquired with two types of 3DE scanners. 18 of these patients were scanned using the fast rotating ultrasound (FRU) transducer [Voormolen et al. 2006], which was connected to a Vingmed Vivid 5 system (GE Vingmed, Horten, Norway).

*FRU
transducer*

The FRU contains a linear phased array transducer that is continuously rotated around its image axis at high speed, up to 480 revolutions per minute (rpm), while acquiring 2D images. A typical data set is generated during 10 seconds at 360 rpm and 100 frames per second (fps). The images of the left ventricle are acquired with the transducer placed in apical position, and its rotation axis more or less aligned with the LV long axis. A single cardiac cycle in general is not sufficient for adequate coverage of the entire 4D space; therefore, multiple consecutive cycles are merged. The cardiac phase for each image is computed offline using detected R-peaks in the ECG [Engelse and Zeelenberg 1979].

preprocessing

The remaining 36 patients, who were referred for stress echo, were scanned using the Philips Sonos 7500 system (Philips Medical Systems, Andover, Massachusetts, USA), equipped with the X4 xMatrix transducer, placed in apical position. For the AAM modeling the FRU data sets are interpolated to Cartesian voxel sets, using a dedicated interpolation method for sparse irregularly sampled data (chapter 4).

Gaussian subsampling was applied to $1/4^{\text{th}}$ of the original resolution for both data types, to reduce speckle in the images.

The data was analyzed using a semi-automatic endocardial border detection method, allowing manual corrections (chapter 2). For the matrix acquisitions, the data was resliced to generate 10 equidistantly sampled long-axis views in all cardiac phases, which is a requirement for the analysis. Full cycle LV endocardial borders were extracted from these analyses and used as the training data sets for the AAM.

*semi-
automatic
analysis*

AAM for the left ventricle in 3DE | 6.3.2

Modeling | 6.3.2.1

Statistical analysis of the left ventricular shapes requires a corresponding point distribution for all the training samples. We define this point correspondence based on key landmark points from the semi-automatic analysis. The points are defined in an anatomical coordinate system of the left ventricle. This cylindrical coordinate system is oriented around the long axis (LAX). Near the apex the surface points are defined in a spherical coordinate system oriented around a center at $3/4^{\text{th}}$ of the LAX (fig. 6.1a). In the cylindrical part, the surface points are sampled equidistantly along the LAX and over the azimuth angle. For the apical part of the surface, sampling is done equidistantly over the elevation and azimuth angle. We chose to define the shape as such, to easily represent the endocardial surface with a regular sampling. This also avoids the need of performing a Delaunay triangulation on the mean shape. The neighboring samples are defined intrinsically in the sampling, which eases the triangulation. The mean shape, with triangulation, is shown in fig. 6.1b. Instead of using the regular definition of the apex, being the point on the surface which is most distant from the mitral valve center, we used a more stable definition. This stable apex is the intersection of endocardial surface with the axis through the center of gravity (CoG) of the upper quarter of the left ventricle (see fig. 6.1c).

shape model

In total, we typically sample at 30 levels from mitral valve to apex, at 30 angles for each level, together with a point for the apex resulting in 901 points for each 3D shape (fig. 6.1b).

We represent the translation, rotation and scaling of the model by 7 pose parameters: 3 for translation, 3 for rotation and 1 for uniform scaling. For the representation of 3D rotation we studied the use of Euler angles and quaternions [Funda and Paul 1988; Horn 1987]. We chose to represent the rotations using quaternions because of the unambiguous representation of 3D orientations, the orthogonality of the representation and the possibility to convert quaternions to rotation matri-

*pose
representation*




Figure 6.1: *a)* A schematic 2-dimensional representation of the regular cylindrical sampling and the spherical sampling in the apical region. *b)* The mean mesh, with triangulation, of a model containing 54 patients. The bottom ring structure is added for texture sampling. *c)* The redefinition of the long axis and the apex using the center of gravity of the apical region

ces back-and-forth without any loss of precision. To avoid any ambiguity in the quaternion representation $q = \{q_0, q_x, q_y, q_z\}$ and to be able to represent the 3D rotation with 3 quaternion parameters $\{q_x, q_y, q_z\}$, we defined that the first, omitted parameter q_0 is always positive (since $q = -q$). Since the norm of the full quaternion $\|q\|$ is 1, the omitted parameter can be recomputed at any time. Scaling is then represented by an independent parameter s .

Procrustes alignment In the shape model we want to model only the biological shape variation of the left ventricle. Therefore the presegmented shapes are aligned using a 3D Procrustes alignment [Goodall 1991; Gower 1975]. In this way, the undesirable absolute position and orientation of the shape are removed from the model. These are a consequence of the acquisition procedure, not of any biological phenomena and are therefore not desirable in the shape model.

texture model The texture sampling has been defined, similarly to the shape model, in the anatomical coordinate system of the left ventricle. We sample the texture radially on the line through the surface points, up to twice the radius of the surface. In this way, the myocardium, part of the right ventricle and a small region outside of the




Figure 6.2: Three orthogonal intersections of the model mean. *From left to right*) Short-axis view, 2-chamber view and 4-chamber view (approximately)

heart is also modeled, to enlarge the lock-in region of the AAM.

Using this anatomical definition we can easily adjust the sampling density to be sparse in regions with little information (blood pool) and dense in important regions (near the endocardium). It also eases the warping of a texture to an arbitrary model shape, which speeds up the computation of the residual vector r . For the warping, we apply a trilinear interpolation defined in barycentric coordinates of tetrahedrons. The texture mean is shown in fig. 6.2. fig. 6.3 shows the three most prominent modes of appearance of the model built on 54 patients.

Since we use a PCA in the modeling of the textures in the AAM, we assume that the texture samples are Gaussian distributed. It is known for ultrasound images that their gray value distribution is non-Gaussian. That is why we apply a non-linear gray value transformation that maps the mean histogram of all the training samples onto a Gaussian distributed histogram with zero mean and unit variance, as introduced by Bosch et al. [2002]. In this procedure, a combined normalized histogram of all the training samples is created. A histogram transformation is defined which transforms this combined histogram into a histogram with a Gaussian gray value distribution. This transformation is then applied to all individual training samples. Subsequently, the training samples are normalized to zero mean and unit variance using the regular gray value normalization used in AAMs.

*gray value
normalization*

Training | 6.3.2.2

In the training procedure we used perturbations δc of {2, 4, 6, 8, 10} mm for translation, and {0.02, 0.04, 0.06, 0.08, 0.1} for the scaling and rotation parameters. A uniform weighting function is used for the weighting of the different perturbations.

*pose and
appearance
parameters*

For training of the parameter update matrix we use perturbations of {0.2, 0.4,




Figure 6.3: 3D renderings illustrating the three most prominent modes of variation from top to bottom with the model mean μ centered




Figure 6.4: Three orthogonal intersections of the column images of the Jacobian matrix J for the three translation parameters. *top to bottom*) translation in x -, y -, and z -direction

0.6, 0.8, 1.0 $\}$ σ for each of the appearance model parameters. These perturbations are uniformly weighted. The perturbation sizes were determined experimentally and correspond to values reported in literature [Cootes and Taylor 2001*b*; Stegmann et al. 2003].

Evaluations of the training procedure, as described in section 6.2.1.3, can be done by visualizing the columns of the Jacobian matrix, warped as textures to the mean shape. Fig. 6.4 shows the resulting images of the columns of the Jacobian matrix J corresponding to the translation parameters. The column images for translation show the high correspondence to the x -, y -, and z -derivatives of the average image, which is as expected.

Regular matching | 6.3.2.3

For the standard AAM matching the parameter update is generated by multiplication of the current residual $r(p)$ with the update matrix U , as in eqn. 6.7. We employ a linear search along the update vector δp using update steps k to minimize the residual $r(p + k\delta p)$. This linear search can be replaced by more sophisticated

variants if desired [Cootes and Kittipanya-ngam 2002]. Consequently, the parameter vector p is updated according to the optimal update step k :

$$p_i = p_{i-1} + k\delta p \quad (6.9)$$

These two steps are repeated until either the model converges or no improvement in r is found (alg. 6.1).

Algorithm 6.1 Regular AAM matching

```

1:  $p_0, r_0 = r(p_0), i = 0$ 
2: repeat
3:    $dp_i = -Ur_i$ 
4:   for all  $k_j \in k$  do
5:      $p_{i+1}(k_j) = p_i + k_j dp_i$ 
6:      $r_{i+1}(k_j) = r(p_{i+1}(k_j))$ 
7:      $dr(k_j) = |r_i|^2 - |r_{i+1}(k_j)|^2$ 
8:   end for
9:   select  $p_{i+1}(k_j)$  and  $r_{i+1}(k_j)$  for largest  $dr(k_j)$ , if  $dr(k_j) \geq 0$ , else break
10:   $i = i + 1$ 
11: until  $i \geq i_{\max}$ 

```

6.3.2.4 | Matching with Jacobian tuning

Cootes and Taylor observed that the assumption that the Jacobian is fixed is unsatisfactory, especially if the image to be segmented is significantly different from the model mean [Cootes and Taylor 2006]. Recently, they have proposed a search strategy that updates the Jacobian matrix during each new evaluation of the residual $r(p)$. The algorithm is closely related to the quasi-Newton methods for solving least square problems without derivatives [Broyden 1965]. The essence of the idea is that during the matching process, we apply variations to the parameters and obtain differences in the residuals. The change in residuals as a result of changes in the parameters provides similar information as is obtained in the regression training process, but in this case the information is highly specific for the case under consideration. Therefore, we would like to tune the standard Jacobian to the current case using the information obtained during the search.

*patient specific
matching*

In short, the method uses a set of constraints on the parameter update at the current iteration i , given all previous parameter estimates $(p_0, \dots, p_i)$ and previous residuals $(r_0, \dots, r_i)$. The Jacobian J_0 from the training phase provides a regularization term for estimating current updates for the Jacobian matrix J . The updated Jacobian matrix J_i is then used to update the appearance parameters. No addi-

tional line search step is required. A summary of the algorithm is given below; for the original derivation, we refer to Cootes and Taylor [2006].

Consider a set of i observations of parameter differences $dp_k = p_k - p_{k-1}$ and residual differences $dr_k = r(p_k) - r(p_{k-1})$, organized in matrices $X = (dp_1 | \dots | dp_i)$ and $R = (dr_1 | \dots | dr_i)$. We set up i linear constraints on each row j_m of J , assuming that a linear update in the parameters generates a linear change in residuals: $X^T j_m = q_m$, where q_m^T is the m^{th} row of R . Using our trained Jacobian J_0 as a regularizer, we can set up a quadratic function of the form $f(j_m) = \alpha |X^T j_m - q_m|^2 + |j_m - j_{0m}|^2$, where α controls the strength of the regularization and j_{0m} is the m^{th} row of J_0 . Differentiating f with respect to j_m and equating to zero leads to an equation for computing a new estimate of J , given the initial estimate from the training set J_0 and all previous parameter updates and residuals:

matching
observations

$$(I + \alpha XX^T)J^T = J_0^T + \alpha XR^T \quad (6.10)$$

where I denotes the identity matrix. eqn. 6.10 can be rewritten into a more efficient version, which is then solved iteratively. Let us define three matrices $A = I + \alpha XX^T$, $B = J_0^T + \alpha XR^T$, and $C = B^T B$. By substituting eqn. 6.10 into eqn. 6.7, it can be shown that the optimal parameter update is given by $dp = Ay$, if y is the solution to the linear equation $Cy = -B^T r$. Instead of calculating A , B , and C using their definitions at every iteration, one can show that these matrices can be updated linearly at the current iteration $i + 1$ using their values at the previous iteration i :

$$A_{i+1} = A_i + \alpha_i dp_i dp_i^T \quad (6.11)$$

$$B_{i+1} = B_i + \alpha_i dr_i dp_i^T \quad (6.12)$$

$$C_{i+1} = C_i + \alpha_i B_i^T dr_i dp_i^T + \alpha_i dp_i dr_i^T B_i + \alpha_i^2 |dr_i|^2 dp_i dp_i^T \quad (6.13)$$

This leads to the Jacobian tuning algorithm for AAM matching, alg. 6.2.

The resulting algorithm has only a series of simple linear operations, and can therefore be added straightforwardly to any existing AAM implementation. Note that the matrices A , B , and C are updated every iteration, regardless of the parameter update. It is usually possible to solve the linear equation in alg. 6.2, l.3 using Cholesky decomposition, as C_i is symmetric and (usually) positive definite. As in Cootes and Taylor [2006], we use $\alpha_i = (\delta + |dp_i|^2)^{-1}$, where δ is small, included to avoid numerical instability after small steps.

Algorithm 6.2 Jacobian tuning AAM matching

```

1: Initialize  $p_0, r_0 = r(p_0), i = 0, A_0 = I, B_0 = J_0, C_0 = B_0^T B_0$ 
2: repeat
3:   Solve  $C_i y = -B_i^T r_i$  for  $y$ 
4:    $dp_i = A_i y$ 
5:    $p_{i+1} = p_i + dp_i$ 
6:    $r_{i+1} = r(p_{i+1})$ 
7:    $dr_i = r_{i+1} - r_i$ 
8:    $z = B_i^T dr_i$ 
9:    $A_{i+1} = A_i + \alpha_i dp_i dp_i^T$ 
10:   $B_{i+1} = B_i + \alpha_i dr_i dp_i^T$ 
11:   $C_{i+1} = C_i + \alpha_i z dp_i^T + \alpha_i dp_i z^T + \alpha_i^2 |dr_i|^2 dp_i dp_i^T$ 
12:  if  $|r_{i+1}|^2 > |r_i|^2$  then
13:     $p_{i+1} = p_i$ 
14:     $r_{i+1} = r_i$ 
15:  end if
16:   $i = i + 1$ 
17: until  $|dp_i|^2 < \epsilon$  or  $i \geq i_{\max}$ 

```

6.4 | Experiments and results

6.4.1 | Model generalization

coupling shape and texture In the proposed AAM a correlation is assumed between the shape and texture of the training samples. That is why the shape and texture model are coupled by an extra PCA on these model parameters. By coupling these models, we might benefit from the correlation between shape and texture, but we also lose some of the variation for the shape and texture that is in the model, since shape and texture are not independent anymore and therefore can only be described together. We investigated the loss of generality by coupling these models.

model truncation Secondly, since the number of training data sets available for training the AAM is limited, we investigated the degree of generalization that is achieved with these data sets. Therefore, we compare the error that is found when projecting a patient's shape or texture on the model in a leave-one-out (L-1-O) evaluation, with the error that is found when projecting the patients on a model with less variation in the shape or texture model, but with the current patient included. This is done by truncating the eigenvector matrix Φ and parameter vector b , such that the re-

maintaining eigenvectors account for the desired amount of variation. The number of eigenmodes t is determined by choosing the t largest eigenvalues $\lambda_0 \dots \lambda_{t-1}$ such that

$$\sum_{i=0}^{t-1} \lambda_i \geq f \sum_{i=0}^{N-1} \lambda_i \quad (6.14)$$

where f is the proportion of the total variation that should be retained and N the total number of eigenvalues. In this way, we can estimate the amount of variation in the L-1-O models, by comparison of the projection errors of the L-1-O models with those of the truncated models. Furthermore, the projection error in L-1-O defines a lower limit on the final segmentation error, when matching to unseen data in a L-1-O evaluation.

Coupled vs. uncoupled models | 6.4.1.1

We evaluated the model generalization of the coupled model versus the independent shape and texture models. The resulting mean projection errors over all L-1-O projections are shown in fig. 6.5. Uncoupling the shape and texture models clearly decreases the projection errors. This is illustrated for the shape model in fig. 6.5. The mean point-to-point (P2P) projection error for the shapes decreases from 1.2 mm for the coupled model to 0.7 mm for the shape model built on 28 patients. A similar decrease in projection can be expected for the texture model. In these experiments we coupled shape and texture models using a weighting factor W_s (eqn. 6.3) that compensates for the difference in units of shape and texture vectors. This weighting factor may also be used to prioritize between shape and texture in the model.

weighting

Comparison with truncated models | 6.4.1.2

We also evaluated the model generalization for the shape and texture models individually and uncoupled, compared to truncated models from all training patients.

The resulting shape projection errors are shown in fig. 6.6. It shows that the shape model generalizes well. For the L-1-0 case with a model containing 53 patients, the model can describe any shape with a mean point-to-point (P2P) distance of 0.41 mm. This corresponds to a projection error of a shape model that contains approximately 97% of the modeled variance (fig. 6.6).

shape

For the generalization of the uncoupled texture models, we compared texture models of raw and normalized textures. For the texture model we express the projection error in mean squared intensity distance (MSD) to the original textures. The texture intensities are normalized to a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 1$. Fig. 6.7 shows that the texture model generalization is

texture




Figure 6.5: Superior generalization levels of uncoupled shape model, w.r.t. the coupled model. Mean projection errors over all the patients on L-1-O models. Point-to-point errors for projection of the training shapes on the shape model of the coupled (black) and uncoupled AAM (red) in L-1-O, for models built on an increasing number of training data sets (x -axis).

weak compared to the shape model, but that the normalized texture model generalizes to a higher level than the raw texture model, comparable to almost 70% and 60% truncated model respectively. Absolute comparison of the projection error is misleading, since the raw texture intensities are also scaled to a distribution with $\mu = 0$ and $\sigma = 1$, while the raw intensities are clearly non-Gaussian and therefore incomparable to the normalized intensities of the normalized texture model.

6.4.2 | Matching evaluation

Since the Jacobian tuning method allowed the training matrix to adapt to the test image, we hypothesized that the method will have a larger capture range. Therefore, we tested the convergence of both methods: the model was initialized at its ideal pose and appearance in the test image, the appearance and pose parameters were then perturbed randomly in a range of several standard deviations, and subsequently the standard AAM and Jacobian tuning method were applied to match the model to the image. The experiments were first performed using a model describing 100% of the shape and texture variation (scenario A). Next, a model was used

matching
scenarios




Figure 6.6: Generalization levels of the shape model as a function of the number of training sets. Mean P2P shape projection errors over all the patients on L-1-O models (redline) vs. the errors obtained by projection on truncated models containing 99%, 98%, 97%, 96%, 95%, 90% and 80% of the total shape variation (dashed lines)

which described only 95% of the shape and 75% of the texture variation (scenario B); in the previous experiments this was shown to be an accurate representation of a leave-one-out situation. Models A and B were built and matched on the same training data. A third scenario (C) was considered, in which models were created in leave-5-out fashion, such that five datasets were reserved for matching and the rest was used for training. This resulted in 11 models (with the last model made by leaving out the remaining four patients).

For all the scenarios we also initialized the models at their mean translation and appearance parameters, to evaluate the matching results when no information about the patient is used, to better approximate a real-life matching situation. The optimal parameters were then found using both matching methods.

For the standard algorithm, update steps $k = [1, 1\frac{1}{2}, 2, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}]$ were used. Matching was terminated if $dr(k_j) < 0$ for all steps k_j . As for the Jacobian tuning algorithm, the matching was allowed to continue until $|dp|^2$ was smaller than $\epsilon = 0.01$. For both methods, the matching was stopped if the mean squares of the residual was smaller than 0.001 (in MSD), or if the maximum of 100 iterations was reached.

stop criteria




Figure 6.7: Generalization levels of the raw and normalized texture models as a function of the number of training sets. Mean texture projection errors (in MSD) over all the patients on L-1-O models (redline) vs. the errors obtained by projection on truncated models containing 90%, 80%, 70%, 60% and 50% of the total texture variation (dashed lines) *a*) For texture models built on the raw textures. *b*) For texture models built on the normalized textures.

Note that MSD units between raw and normalized texture models cannot be directly compared (see section 6.4.1.2)

6.4.2.1 | Perturbation from ideal parameters

scenario A Point-to-point errors between the matching results and the manually drawn contours were calculated. With a model describing 100% of the shape and texture variation (scenario A), very low matching errors could be expected. For this experiment, a matching was considered converged if the point-to-point (P2P) error, averaged over the contour, was lower than 1 mm (the largest voxel size). The results revealed that the Jacobian tuning algorithm was superior to the standard algorithm (see fig. 6.8). In this case, 14.2% (69 out of $54 \times 9 = 486$) did not converge using the standard matching algorithm, whereas the Jacobian tuning algorithm achieved a 100% convergence rate. Most outliers occurred because the standard algorithm was not able to find an update for all steps k_j during the first iteration, such that the residual was lower than the residual at initialization.

scenario B & C As for the truncated model (scenario B) and the leave-5-out models (scenario C), a lower spread in errors and much higher accuracy was observed if using the




Figure 6.8: Superior matching convergence of the Jacobian tuning algorithm over the standard AAM matching. Perturbation experiments using a model with 100% shape and texture variation, initialized with perturbations of 0.1, 0.3, 0.5, 0.7, 1.0, 1.5, 2.0, 2.5 and 3.0 σ . Boxes indicate 25% and 75% percentiles; whiskers extend to 10% and 90%; red dots indicate cases with a final matching error above 1 mm

Jacobian tuning algorithm, especially when large perturbations were applied (fig. 6.9). Similar results obtained for scenario C are shown in fig. 6.10.

The difference in the results of the Jacobian tuning matching between scenario A and B show the impact of the truncation of the shape and especially the texture model. Just a truncation of the shape model would presumably yield P2P errors that approach those of the projection experiments (fig. 6.6), close to 0.4 mm. However, in the matchings on the truncated model (scenario B) we found much larger errors, around 1.8 mm. This increase can be primarily attributed to the weak truncated texture model.

truncating the models

Initialization at mean parameters | 6.4.2.2

The matching results for initialization at mean translation and mean appearance parameters are given in table 6.1. Significantly lower errors P2P and P2S errors were obtained with the Jacobian tuning algorithm. A segmentation example is shown in




Figure 6.9: Superior matching performance of the Jacobian tuning algorithm (red) over standard AAM matching, compared to initial errors. Perturbation experiments using a model with 95% shape and 75% texture variation, initialized with perturbations of 0.1, 0.3, 0.5, 0.7, 1.0, 1.5, 2.0, 2.5 and 3.0 σ (slightly pulled apart for clearer visualization)

fig. 6.11.

As expected from the previously described perturbation experiments, matching errors again increase from scenario A to B and C. This is due to the diminished level of generalization of the truncated and L-1-O models. Again, especially the weakness of the texture model seems to contribute most to the increased matching errors.

computation time The computation time required for the AAM matching depends mostly on the number of (allowed) iterations. While there is still much room for optimizations, an iteration takes 1-2 seconds on regular desktop PC (2-3 GHz processor). Compared to the regular AAM matching, the Jacobian tuning takes twice as much time per iteration, but needs less iterations to come to convergence. Thus Jacobian tuning can have a similar computation time if compared to the regular matching.




Figure 6.10: Superior matching performance of the Jacobian tuning algorithm (red) over standard AAM matching, compared to initial errors. Perturbation experiments using a leave-5-out model, initialized with perturbations of 0.1, 0.3, 0.5, 0.7, 1.0, 1.5, 2.0, 2.5 and 3.0 σ (slightly pulled apart for clearer visualization)

Discussion | 6.5

General conclusions | 6.5.1

We have successfully developed a fully automatic method for segmentation of the left ventricle in 3D echocardiography based on active appearance models. The method has shown to provide good segmentation within a set of constraints. We have explored and compared variations and extensions of the classical AAM approach, and have charted the current boundaries of applicability in our problem domain. We have shown the generalization capabilities of the derived statistical models, discussed the importance of a proper choice of modeling, training and matching parameters, and showed the considerable added value of enhancements such as the Jacobian tuning matching approach. The required amount of computation time was acceptable for practical applications. Although a number of issues still needs to be tackled, AAMs constitute a very promising approach for automated




Figure 6.11: Appearance patches and 3D segmentations results using the regular AAM and the Jacobian tuning algorithms. For all segmentations a short-axis slice and two long-axis slices are shown. For the two segmentation approaches the detected model appearances are shown with contours and also a checkerboard image that combines the detected appearance with the original image and manual contour (redline). On the right the detected surface (solid) is shown together with the manual surface (mesh). In this particular case, the manual gold standard is very different from the mean appearance (fig. 6.2). The regular AAM matching has trouble finding the correct segmentation, as opposed to the Jacobian tuning method

Table 6.1: Mean $\pm$ standard deviation of P2P matching errors when initialized at mean translation and mean appearance parameters. Also shown are the root-mean-squares intensity error (RMS), in unnormalized intensity units (range [0 255]).

*indicates that Jacobian tuning has statistically significantly better results than standard AAM (paired t -test, $p < 0.05$, $N = 54$)

Measure	Match	Scenario		
		A	B	C
P2P (mm)	Initial	7.5 ± 2.6	7.5 ± 2.6	7.5 ± 2.7
	Regular AAM	2.3 ± 1.1	3.6 ± 1.7	4.4 ± 1.7
	Jacobian tuning	$0.06 \pm 0.03^*$	$2.9 \pm 2.0^*$	$3.9 \pm 2.0^*$
P2S (mm)	Initial	4.4 ± 1.2	4.4 ± 1.2	4.4 ± 1.2
	Regular AAM	2.0 ± 0.6	2.6 ± 0.5	3.0 ± 0.7
	Jacobian tuning	$0.06 \pm 0.03^*$	$2.2 \pm 0.7^*$	$2.8 \pm 1.0^*$
RMS intensity	Initial	5.2 ± 2.5	5.2 ± 2.5	5.2 ± 2.6
	Regular AAM	3.4 ± 1.7	3.6 ± 1.5	3.9 ± 1.7
	Jacobian tuning	$2.0 \pm 0.8^*$	$3.5 \pm 1.6^*$	3.8 ± 1.7

segmentation of the LV in 3DE that yearns for further substantiation.

We will discuss the advantages of matching using Jacobian tuning with respect to the classical matching approach, the current limitations, related work and conclude with recommendations for further research.

Regular matching vs. Jacobian tuning | 6.5.2

This study demonstrates, among others, the effectiveness of the new Jacobian tuning matching approach in AAM segmentation of the left ventricle in real-time 3D ultrasound images. We showed that the Jacobian tuning algorithm has a larger capture range and higher accuracy than the standard matching algorithm.

It is interesting to see that the outliers in fig. 6.8 are all located above approximately 4mm, suggesting that, below this threshold, it is possible to find the optimal appearance parameters using the standard algorithm. The Jacobian tuning method is much more robust because of its larger capture range, obtaining a 100% success rate for perturbations up to 3 standard deviations (σ) from the ideal parameters.

outliers

Another interesting observation is the lower bound of 2 mm P2P error for the

capture range truncated and leave-5-out model, which can be achieved for perturbations up to 3σ using the Jacobian tuning method, whereas the standard algorithm starts to fail around 2σ perturbation from ideal parameters (fig. 6.9 and 6.10). Of course, these error bounds and perturbation limits are dependent on the amount of variation captured in the model. However, it is clear that the Jacobian tuning method has a much larger capture range than the standard AAM algorithm. This may have significant consequences in matching models to images acquired with different machine settings and transducer equipment. For example, it would be worth experimenting with a model built with Philips data and matched to FRU data. This is a subject of further investigation.

6.5.3 | AAM matching strategies

Other AAM search algorithms have been reported in the literature, which implement updates to the Jacobian matrix. For example, Batur and Hayes [2005] proposed an algorithm which uses linear updates for the gradient matrix. Their approach is different to this one in the sense that the current parameters of the texture model are used to update the appearance parameters. This is combined with a line search similar to the one in section 6.3.2.3, and matching is stopped if no better residuals are found. This is different in our approach, where the Jacobian can be updated infinitely if desired. The Jacobian tuning method is closely related to the quasi-Newton method for solving least-squares problems without derivatives proposed by Broyden [1965]. More sophisticated approaches were proposed by Xu [1990]. These types of algorithms merit further research and comparison.

6.5.4 | Current bounds of applicability

6.5.4.1 | Rotation matching

small capture range No results for perturbations in rotation are reported in this study. In all our evaluations we have seen that adding the rotation parameters to the optimization process raises problems in the optimization. The capture ranges in perturbation experiments including rotation parameters are relatively small compared to capture ranges of other (pose) parameters and the matching accuracy degrades quickly outside these capture ranges for rotation parameters. These limitations could be attributed to several causes.

Firstly, we can attribute these problems to the high degree of rotational symmetry of the left ventricle. Both the shape and the texture show only small differences

with respect to rotation around the LV long axis, which makes the objective metric relatively insensitive to these rotations. Therefore, also the training and matching of the AAM show poor performance in characterizing and solving these rotations respectively. *rotational symmetry*

Secondly, the detection of the rotation parameters is likely hampered by typical ultrasound specific image characteristics. Most acquisitions suffer from drop outs as a result of rib shadowing. Current 3D ultrasound equipment employs a footprint that is not small enough to image in between the ribs without resulting in shadows in parts of the image, especially the left ventricular lateral wall. Also, prominent image artifacts are common in the near-field, making it difficult to locate the LV apex precisely. Together, these typical image artifacts challenge the detection of the orientation of the LV. *typical image artifacts*

Furthermore, the myocardial texture may vary considerably throughout the left ventricle, due to misalignment of the acquisition axis with the left ventricular long axis. This is a typical ultrasound specific feature, caused by the angle of incidence of the ultrasound beam to the myocardium. This hampers the detection of the correct rotation with respect to the LAX since this typical change in texture does not correspond to any biological variation in shape. *angle of incidence*

Finally, the inherent nonlinearity of rotation representations may play an important role. From experimentation with Euler-angle representation vs. quaternion representations, we found that quaternions are better behaved with respect to numerical stability, but still the implementation choices might limit the lock-in range for rotations. The optimal representation of rotation angles can be a field for further investigations.

Nonlinear gray value normalization | 6.5.4.2

Another limitation of the current implementation of the AAM for 3DE is the nonlinear gray value normalization. AAM modeling expects a Gaussian distribution of variability, both for point distributions and textures. Ultrasound gray value distributions are known to be non-Gaussian. The nonlinear gray value normalization compensates for the non-Gaussian distribution of the gray values in ultrasound images. It is trained to transform the (linearly) normalized ultrasound histogram into a histogram that approximates a Gaussian distribution. Since this normalization is learned from the set of sample textures, it is limited to normalization of textures that have a comparable normalized histogram to the sample textures. Especially in cases where the AAM is far away from the optimal pose, the sampled texture histogram may be considerably different; e.g. if the scaling factor is too small, the texture contains mainly blood values. *learned normalization*

Alternatives to the gray value normalization scheme merit further research. Im- *preprocessing*

provements may be found in preprocessing steps that try to enhance the textures directly, by emphasizing desired features and masking out undesired ones, such as ultrasound speckle. This is discussed in section 6.5.5.4.

texture modeling Other improvements may be found in the modeling of the textures. Larsen et al. [2007] describe the intensity information using wavelets and wedgelets, resulting in a more compact representation of the texture model. This might benefit both the distribution of the gray values and the generalization of the texture model. The generalization of the texture model may also be improved by using a cluster-aware AAM that employs clustering of the texture space to divide the non-Gaussian distributed texture model in multiple Gaussian submodels that describe the entire texture variation [Stegmann et al. 2005].

6.5.4.3 | Additional constraints

technical imaging limitations Current 3D ultrasound imaging systems are not capable of truly real-time imaging of the whole heart. The speed of sound in human tissue is the limiting factor that determines the maximum number of beam that can be sequentially acquired per second. Therefore, images of multiple consecutive cardiac cycles are merged to obtain a high resolution image of the full left ventricle in 3D over the entire cardiac cycle. This multi-beat fusion often generates typical image artifacts, as a result of inter-beat variation, respiratory motion, patient's motion or transducer motion during the acquisition. These artifacts challenge the feasibility of locating the true endocardial border in these images, even by experts.

limited to model variation AAMs are based on a statistical model of the shape and texture of a certain population of example images with their contours. Training and matching of the AAM relies on correspondence between model parameter variations and the change in the residual image. This relation is learned from the example data sets. Therefore, AAMs are limited to detection of contours which can be described by the statistical variation in the shape model, in images that can be reasonably approximated by the combined shape and texture model. Possible pathological cases, or biological subsets that are not represented in the training population will be detected with limited precision.

manual contours as gold standard Furthermore, there may be imperfections in the manually segmented example data sets with which the models are built. These imperfections degrade the point correspondence of the shape model and also influence the quality of the texture model. Specifically, the alignment of the rotation with respect to the long axis is problematic. Often the right ventricle's attachment point and the aortic valve are hardly visible in the images; this hampers correct determination of this orientation.

Recommendations | 6.5.5**Temporal domain | 6.5.5.1**

In our approach, we have focussed on segmentation of the left ventricle still 3D echocardiography, while much information about the myocardium is in the dynamics of the heart. Therefore, analysis of the temporal domain could greatly enhance also the detection in the single cardiac phases.

Extensions can be made, in parallel to single phase detection, to both the shape and the texture modeling of the AAM. Frangi et al. review various time-varying modeling techniques in cardiac applications [Frangi et al. 2001].

Even more information might be taken from temporal analysis of the texture. This may reveal differences between drop outs and shadowing on one hand, and the lower intensities of the blood pool on the other hand. This distinction is hardly extractable from single phase analyses. Static features in general, are not expected to provide much information about the functional parameters. Integration of temporal features with a dynamic shape model in echocardiography have been proposed by several groups. Comaniciu, Zhou et al. present a tracking framework with an adaptive shape model, where information about measurement uncertainty in the feature detection is combined with shape information using a Kalman filter approach. The framework has been shown to work well in 2D echocardiography [Comaniciu et al. 2004; Zhou et al. 2005]. A similar approach had been proposed by Jacob et al. [2002].

Hybrid matching | 6.5.5.2

The standard AAM matching, as proposed by Cootes and Taylor [2006] solves the update problem by using a fixed update matrix that is estimated at the optimal position. This type of matching assumes that the current model state is close to the optimum and is thus limited to a so-called lock-in range. Also, it only allows updates to the model parameters and is therefore bound to global model updates and to the variation in the model. This usually leads to a globally correct segmentation with minor mismatches. However, local mismatches may even cause the global segmentation to deteriorate.

Jacobian tuning has been shown to have a larger lock-in range and to result in a better fit of the model to the data. However, it is still limited to the variation in the model. Further improvement in local matching of the AAM may be found in a hybrid matching approach, like a combination with (multidimensional) dynamic programming.

6.5.5.3 | Multidimensional dynamic programming

The strength of the regular AAM matching in quickly finding a global solution also has a drawback. Finding a global solution may sometimes result in locally suboptimal segmentations. This may especially occur in pathological cases, which may not be covered by the samples in the model training. Therefore, regular AAM matching could be extended with multidimensional dynamic programming [Üzümcü et al. 2006] (MDP, an extension of dynamic programming (DP) [Amini et al. 1990; Bellmann 1965; Sonka et al. 1999]), to be able to fit the model's appearance more locally to these cases, without harming the global optimization. MDP can be used in this case to locally match the texture, synthesized by the regular matching, to the underlying unseen image. Since MDP is a path search strategy and we aim at detecting an optimal surface in 3D, MDP could be applied in collinear planes (slices) of the 3D volume, intersecting the long axis. This yields a set of 2D contours that constitute the 3D endocardial surface. From this new surface, a pose update can be computed and the shape can be projected onto the shape model to update the appearance model. The use of MDP (instead of regular dynamic programming) allows free movement of the surface points on the plane (or even out of plane) in the MDP update step. This could make the model more easily adapt to rotation errors for example.

We previously applied DP search in a semi-automatic segmentation approach for 3D echocardiography [van Stralen et al. 2005]. Oost et al. [2006] have successfully applied a hybrid approach to multi-phase X-ray angiograms, Hansegård et al. [2007a] extended this work to multi-view, multi-phase echocardiography. Therefore, an extension of the 3D AAM with (multidimensional) dynamic programming might be a promising subject for further research.

6.5.5.4 | Preprocessing

Echocardiographic images have a characteristic granular appearance, referred to as speckle [Wells and Halliwell 1981]. Speckle is often seen as undesirable in static (single frame) ultrasound image analysis and therefore we chose to simply reduce the speckle by Gaussian smoothing. More advanced preprocessing techniques have been proposed for speckle reduction. Among them are wavelet-based methods [Achim et al. 2001; Xiao et al. 2004] and anisotropic diffusion methods [Tauber et al. 2004; Yu and Acton 2002]. Nillesen et al. apply adaptive filtering using image statistics as a preprocessing step for automated image segmentation [Nillesen et al. 2007; Nillesen et al. 2008]. This adaptive filter enhances blood-to-tissue contrast and may therefore be used as a preprocessing step in the AAM modeling and matching.

Methods based on anisotropic diffusion may be easily incorporated in the re-

construction process of irregularly distributed ultrasound data, as proposed in chapter 4.

Cardiac phase | 6.5.5.5

A single phase segmentation method provides an important base for detection of left ventricular functional parameters, e.g. by combining it with a tracking approach. In such a framework, one is free to choose any cardiac phase for the initial segmentation.

We studied automated segmentation of the left ventricle in the end-diastolic (ED) phase. This cardiac phase allows clear visualization of the endocardial wall since it can be easily identified by the R-peak in the ECG and it does not suffer from much wall motion. A drawback of the use of the ED phase for LV segmentation is that it is the phase where the LV reaches its largest volume. Therefore, parts of the LV (especially the apex) might not be captured in the field of view or might be shadowed by the ribs or lungs. Also, in ED the endocardial wall is less clearly defined due to possible visibility of individual trabeculae, forming an irregular undulated surface [Mannaerts et al. 2003].

*detection in
end diastole*

An obvious alternative to detection in ED would be single phase detection in ES. The advantages would be that the ventricle is more likely to be fully captured in the image volume and shadowing artifacts might be less prominent. Also, we found in chapter 5 that the acquisition axis is closer to the LV long axis, which eases the initialization. Also, the endocardial border definition might be more clearly defined [Mannaerts et al. 2003]. Difficulties of applying detection in ES might be proper detection of the ES phase from the ECG and the fact that LV geometry in ES may vary more widely than in ED, as a result of various pathological contraction patterns (e.g. the occurrence of 'kissing ventricles').

*detection in
end systole*

Uncoupled AAM | 6.5.5.6

The generalization experiments in section 6.4.1 show that with the current size of the training set, the texture model only generalizes moderately. In cases where only a small training set is available, compared to the intrinsic dimensionality of the data, the detection can be improved by decoupling the shape and texture models. For example, Stegmann and Pedersen [2005] have used these uncoupled appearance models in segmentation of cardiac MR.

A drawback of decoupling the shape and texture models is that correlations between shape and texture modes are no longer modeled. Furthermore, the dimensionality of the model increases. This may hamper the matching process.

Theoretically, uncoupled models are able to describe unseen patients more precisely than coupled models. However, this difference can only be appreciated if the

matching proceeds to a high correspondence with the target image. Only then, the difference between the coupled and uncoupled model can be seen. As long as a good global match is not yet achieved (in pose, shape and texture), the uncoupled model will not be able to tune in great detail and match to the unseen data optimally.

6.5.5.7 | Improved/manual initialization

AAMs follow an iterative matching approach in search for a global optimum. This optimization is not guaranteed to find this global optimum. The quality of the initialization of the AAM influences the final segmentation result. Thus, the better the initialization, the better the final segmentation. There are several options to improve the initial model and to process the given initial data.

automated initialization The most common way of initializing an AAM is by initialization on the mean position and orientation and the mean model parameters. In chapter 5 we presented a method to improve the initial transformation parameters automatically using a Hough transform and multidimensional dynamic programming. Zagrodsky et al. [2005] use a computationally costly atlas registration as initialization.

manual initialization Most automated analyses techniques however, are initialized manually by a human observer. Initialization is done either by explicitly annotating the apex, a number of points on the endocardial border or the mitral valve [Corsi et al. 2002; van Stralen et al. 2005], or by indicating the LV position and dimensions by annotations [Angelini et al. 2001; Gérard et al. 2002; Kühl et al. 2004].

Cootes and Taylor [2001a] also describe a very elegant way to constrain the AAM, by including priors on the point positions. In this way, manual or automated feature points can be used to guide the AAM search. Hansegård et al. [2007a] followed this approach to combine the regular AAM optimization with dynamic programming on multi-plane echocardiography.

6.6 | Conclusions

We investigated the use of AAMs for automated segmentation of the left ventricle in 3D echocardiography. The shape and texture model were based on anatomical reference points, aiming at optimal sampling of the left ventricle image region.

We evaluated the degree of generalization of an appearance model based on 54 patients and found that the variation in shape can be sufficiently described by these 54 samples. L-1-O projection errors on such a model are comparable to projection

errors on a truncated model, containing 97% of the total shape variation. Because of the non-Gaussian distribution of ultrasound intensities, we compared generalization levels of raw and nonlinearly normalized texture models. The normalized texture model showed better generalization than the raw texture model. In contrast to the shape model, the generalization level of the texture models was low: projection errors are comparable to a those of a model that is truncated to 70% of the texture variation. This poses a challenge on the AAM segmentation. Either increasing the training set, or preprocessing the texture to get a more efficient representation of the texture variation should be employed to overcome this limitation.

We evaluated matching of the AAM in different scenarios, including variations in initialization, generalizations levels and matching strategies. Comparison of matching results using full and truncated models showed the impact of the weak generalization of the texture model. This is also reflected in the leave-one-out matching experiments.

The Jacobian tuning algorithm for AAM matching produces quantitatively better segmentations than the regular AAM matching, showing larger capture ranges and a higher accuracy. These experiments demonstrate the effectiveness of an adaptive update matrix during model matching. Given the large variability in ultrasound image appearance, the Jacobian-tuning algorithm has great potential in improving the standard AAM segmentation.

The main source of errors in matching the AAM are in finding correct rotation parameters. We excluded the evaluation of rotation matching from our analyses. Apart from this limitation, which can be circumvented by a proper initialization, 3D AAMs have shown to provide a valuable global segmentation of the left ventricle. Further research on combining the promising AAM approach using local matching strategies, temporal image analysis and improved initialization procedures is necessary to achieve a clinically acceptable automated segmentation approach.

Acknowledgments |

The sonographers and clinicians of the Erasmus MC Thoraxcenter, in particular Boudewijn Krenning, Wim Vletter, Jackie McGhie, Attila Nemes, Osama Soliman, Marcel Geleijnse and Folkert ten Cate, are gratefully acknowledged for acquiring the images.

Discussion and conclusions

7

7.1 | Research objective

objective

Assessment of left ventricular function is important in clinical decision making, evaluation of therapeutic effect and determination of prognosis. Automating this assessment may improve the diagnosis and treatment, and reduce its costs through reduction in time and labour needed for the analysis. It may also reduce interobserver and interinstitutional variability and improve the reproducibility.

We aimed at automating the assessment of left ventricular function using 3D echocardiography, especially by employing the cost-effective design of the fast rotating ultrasound (FRU) transducer.

7.2 | Contributions

semi-automatic quantification

Most important in evaluation of LV function is the estimation of the LV volume over the full cardiac cycle. We have shown that this full cycle analysis can be achieved semi-automatically using a method based on dynamic programming and pattern matching with good correspondence to our golden standard, without most of the geometrical assumptions that are made in analysis of 2D echocardiography.

automated tracking

The main source of errors in the full cycle analysis lay in the weak temporal model for mitral valve motion. We studied automatic analysis of this motion by tracking the mitral valve hinge point using multidimensional dynamic programming and apodized block matching in VCR recorded 2D echocardiograms. It shows that this tracking can be reliably done over the full cardiac cycle with an accuracy and precision that are close to those of the human observer. Its versatility allows it to be tuned to either low level image structures such as speckle, or more coarse anatomical structures such as the septal mitral valve hinge point, and it has the conceptual ability to be applied to 3D tracking. The robustness of the tracking in low quality and relatively low frame rate interlaced VCR recorded images is promising for future application in 3D echocardiography using the FRU transducer, which has better image quality but an effectively lower temporal resolution in 3D.

image reconstruction

The FRU transducer features high quality 2D frames, compared to competing 3D matrix transducers. However, reconstruction of the 3D image sequences for visual inspection of the LV geometry and function, is challenging because of the 2D nature of the acquisition and the irregular and sparse distribution of the 2D frames in 3D+T space. We proposed a reconstruction method based on normalized convolution to improve on regular trilinear interpolation using temporal binning. Our

new method deals with the 4D (3D+T) nature of the data and visually improves the temporal delineation of anatomical structures. The interpolation framework is flexible and has room for improvements in development of adaptive filters that optimize the interpolation spatially and temporally. Adjustment of the kernel dimensions to the sparseness of the data can improve the interpolation both in dense and sparse regions.

The interpolation of the FRU data to volumetric sequences allows application of generic 3D+T analysis methods. We developed a method for automatic detection of the LV long axis and the mitral valve plane in 3D echocardiography. The detection of these skeletal features of the LV can be done with good accuracy and precision, promising applicability as an initialization step for fully automatic model based segmentation approaches. The method showed good agreement with manual annotations and may therefore save valuable time and labour in such applications, removing also interobserver variability.

*automatic
initialization*

Finally, fully automatic segmentation of the left ventricle in 3D echocardiography has been investigated using active appearance models. In this work, the value of the dedicated gray value normalization is reflected in the superior generalization of the normalized texture model. Matching optimization using Jacobian tuning substantially improves the accuracy of the detection and especially the lock-in range. These controlled matching experiments in in-vivo data show the feasibility of LV volume quantification using AAMs in 3D echocardiography and its future promise.

*model based
segmentation*

Discussion | 7.3

In general, we have succeeded in developing several techniques that each bring analysis of 3D echocardiography further to its ultimate goals: making it more reliable, less observer dependent, cheaper and less labour intensive. Certainly, in near future echocardiography will improve technically, resulting in better image quality and higher frame rates available at a more affordable price, allowing better automated analysis. In this work, we have shown that automated analysis based on 3D echocardiographic data acquired with the cost effective FRU transducer can be competitive with commercial matrix transducer systems. We have proposed generic techniques for automated reconstruction, tracking, feature detection and segmentation that reduce interobserver variability and time of analysis, making analysis of 3D echocardiography more reproducible and cost effective.

Most of these techniques are applicable or extendible to related ultrasound mo-

general applicability dalities. The automated segmentation and automated initialization procedures (chapter 2, 5 and 6) might directly be employed in 3D transesophageal echocardiography (TEE). Segmentation using AAMs might also be extended to simultaneous segmentation of multiple cardiac chambers as far as those can be visualized properly with 3D echocardiography. The 3D+T interpolation based on normalized convolution (chapter 4) might also be applicable to reconstruction of freehand 3D ultrasound in various clinical applications, e.g. in intraoperative neuro imaging. Future applications of the initialization procedure in chapter 5, might also be found in centerline detection in 3D vascular imaging, e.g. of the carotid arteries. Automated tracking based on multidimensional dynamic programming (chapter 3) can be of great value in most echocardiographic modalities in 2D and 3D. This feature-based tracking needs only a low frame rate and gross image structure, allowing it to complement or enhance the more common speckle-based tracking, which demands high frame rates and high inter-frame correspondence.

image quality Currently, image quality is a major limitation in (automated) analysis of 3D echocardiography. Current transducers suffer from a relatively large footprint, often resulting in rib shadowing. Besides, shadows resulting from the lungs are an issue in apical acquisitions. These limitations restrict apical 3D echocardiography to imaging of the left ventricle only, especially in stress protocols, and even in the left ventricle major shadowing artifacts are often present. Transducer sensitivity and bandwidth are also compromised in current 3D transducers, resulting in sub-optimal harmonic capabilities and higher levels of noise and artifacts than in 2D images. Furthermore, image resolution is limited, especially in the far field regions of the image. In FRU data, sparseness is also present due to the irregular sampling of the 3D+T space. The typical tradeoff between spatial and temporal resolution (volume frame rate) still poses a major limitation and keeps both below the desired level. This may soon be alleviated by massively parallel beam forming, but currently the only remedy is multi-beat fusion. This multi-beat fusion in similar near real-time solutions leads to stitching artifacts as a result of inter-beat variability and patient's, transducer or respiratory motion. These artifacts challenge the feasibility of locating the true endocardial border in these images, even by experts.

technical limitations Enhancement of the endocardial border in 3D echocardiography can be achieved through the use of contrast agents for LV opacification. Although this may improve the visual distinction between blood and myocardium, it might not directly benefit automated analysis techniques. Contrast agents may render spurious edges in static images, that actually are a result of contrast swirls, complicating the automated detection. Besides, contrast agent concentration may vary during the acquisition and high concentrations may yield high attenuation, both resulting in variable contrast-to-tissue ratios. Automated techniques that combine a model-

contrast agents

based approach (chapter 6) with an adapted edge detector may be most promising in analysis of these images.

Recommendations for future work | 7.4

Widespread use of 3D echocardiography will require robust and reliable analysis tools to ensure its value in diagnosis and clinical decision making. Improvements in these automated analysis techniques may be found in further investigation of the following directions.

Texture analysis | 7.4.1

As we observed in chapter 6, myocardial texture may differ widely across the position within a single acquisition as a result of rib or lung shadowing, the angle of incidence of the ultrasound beam, attenuation, anisotropy, near-field clutter, noise etc. In our case, this resulted in poor generalization of the texture model. Future work should be directed towards identification of these typical artifacts, such that a reliability measure can be taken into account for each image region. Much can be gained by temporal analysis of the data, revealing functional information from seemingly noisy speckle patterns. Also preprocessing, optionally integrated into the reconstruction process, for example using spatiotemporal adaptive filtering, might improve the essential signal-to-noise ratio in the relatively low quality ultrasound images.

Model based segmentation | 7.4.2

Despite any future improvements in ultrasound technology, 3D echocardiographic imaging will remain a challenge and associated with previously mentioned typical artifacts. Therefore, segmentation of the left ventricle in 3D echocardiography requires knowledge about the heart as an organ; its anatomy, physiology, its dynamics. Typical artifacts will mislead automated detection techniques, which should overcome partly missing data, false edges as a result of the reconstruction or shadowing and so on. Without any knowledge about the target object or organ, these artifacts cannot be overcome. As shown in chapter 6, statistical models, such as

AAMs, are a very elegant way of incorporating all this knowledge about anatomy, dynamics and appearance together with typical artifacts. Increasing the training populations will further improve the generalization capabilities of these models. Also, more efficient representation of textures may be found using e.g. wavelets or speckle reduction techniques, and the texture model may be better preserved in the AAM by decoupling from the shape model.

7.4.3 | Integration of presented techniques

Above all, previously presented techniques in this thesis should be integrated. At first, initialization of a model based approach might be improved using automated feature detection techniques as in chapter 5. Also, such a global segmentation technique may be extended with a spatially local refinement technique such as (multidimensional) dynamic programming, to relax the constraints from the statistical model, and improve matching to local image features. Single-phase detection using AAMs could be a basis for tracking approaches using for example MDP, providing the desired functional information such as ejection fraction, stroke volume etc., while avoiding complex 4D modeling.

7.5 | Conclusions

In conclusion, we have shown that we succeeded in developing methods that make analysis of 3D echocardiography more reliable and less user demanding. We improved on labour intensiveness and reliability of semi-automatic endocardial border detection. Automated tracking of the mitral valve can be done effectively, with high accuracy and precision in 2D echocardiography. Furthermore, we improved the general applicability of 3D echocardiography using the FRU transducer by achieving better reconstructions from the 3D+T data. This allowed us to develop methods that were not limited to data from the FRU transducer, but also applicable to data from commercially available matrix transducer systems. This led us to robust techniques for detection of the left ventricular long axis and the mitral valve plane. And finally, we showed the promise of fully automatic segmentation of the left ventricle using active appearance models.

The presented techniques may be of great value in related fields, such as reconstruction and analysis of transesophageal 3D echocardiography, freehand 3D echocardiography and 3D carotid imaging. Besides, advanced model-based segmenta-

tion using AAMs may benefit other modalities such as segmentation in cardiac MR and CT and may also improve from developments in those fields.

As discussed in the previous section, there are many opportunities to improve automated analysis of 3D echocardiography. Further development and integration of the presented techniques on reconstruction, feature detection, segmentation and tracking, and improvements in 3D ultrasound technology increasing the spatial and temporal resolution, may soon lead to reliable fully automatic solutions for left ventricular functional analysis based on our work.

Summary

8

The left ventricle of the heart is responsible for pumping oxygen saturated blood through the whole body. Assessment of left ventricular function is therefore very important in clinical decision making, evaluation of therapeutic effects and determination of prognosis. In the last decade, 3-dimensional (3D) echocardiography has evolved into a valuable imaging technique for evaluation of left ventricular function. This thesis, as posed in chapter 1, aims at automating the analysis of 3D echocardiography to reduce interobserver variability and analysis time, and to improve reproducibility of functional assessment of the left ventricle. In this way, diagnosis and treatment of cardiovascular disease may be improved and its costs may be reduced. In this chapter, an overview of the basics in ultrasound imaging is given. Current 3D echocardiography, including the fast rotating ultrasound (FRU) transducer and commercial matrix transducers, is put into historical perspective. Also, major issues in digital image analysis are discussed, with a focus on automated analysis of echocardiography.

Chapter 2 describes a method for endocardial border detection for reliable assessment of full cycle left ventricular volume, requiring minimal user interaction. In this chapter analysis focuses on data acquired with the FRU transducer. The method is based on detection of the endocardial border in the 2-dimensional (2D) slices using dynamic programming and pattern matching. 3D plus time (3D+T) patient specific shape and texture models were derived from four manually drawn contours for each patient, describing the approximate shape and texture of the endocardial border for this patient. Knowledge in these models was optimally combined using dynamic programming and pattern matching to detect the contours in all the required image slices throughout the entire cardiac cycle, such that the 3D endocardial surface could be reconstructed and volumes could be estimated over the full cardiac cycle. The performance of this semi-automatic technique was evaluated on 10 patients and compared to gold standard MRI determined volumes. A good correspondence to MRI volumes was shown. However, the main advantage of this method, being the possibility to manually correct automatically detected contours, was also evaluated and showed an improved correlation to MRI volumes. Also, a low interobserver variability was found for the semi-automatic method.

Chapter 3 discusses one of the main disadvantages of the semi-automatic border detection: the poor model for mitral valve displacement throughout the cardiac cycle. Hence, an automated technique for tracking the mitral valve hinge points (MVHPs) was investigated. To enforce continuous and cyclic motion of the MVHPs while allowing displacements in horizontal and vertical direction, tracking was done using multidimensional programming with apodized block matching. The method was tested on VCR recorded 2D echocardiograms of clinical quality, to evaluate the robustness of the tracking, and compared to common tracking techniques. High accuracy and precision was found, comparable to interobserver variability of man-

*research
objective*

*FRU
transducer*

*semi-
automatic full
cycle analysis*

*mitral valve
tracking*

ual tracking, showing the value of this approach and promise for application in 3D echocardiography.

In chapter 4 reconstruction of the 3D+T data acquired with the FRU transducer is studied. The typical challenges in this reconstruction process are discussed. A solution that deals with irregularity and the sparseness of the data is presented, based on normalized convolution. The presented approach is demonstrated on simulated data and tested on in-vivo FRU data. The reconstruction shows advantages over the standard reconstruction technique, reducing multi-beat fusion artifacts and spurious edges in sparse regions. The method is generic and can be applied to a wide range of irregular 3D ultrasound data, such as freehand 3D ultrasound or mechanical 3D transesophageal echocardiography.

*3D+T image
reconstruction*

Chapter 5 deals with a common problem in model-based segmentation approaches: the model initialization. To this extent, a skeletal feature detection approach is presented that fully automatically determines the left ventricular long axis and the position of the mitral valve plane. These two features may be very useful in initialization of model-based approaches. The presented method reliably detects the long axis using dynamic programming and a Hough transform for circles, based on detection of the LV center in slices perpendicular to the estimated long axis and tracking of these centers over the full cardiac cycle. Using the detected long axis, also a reliable estimate of the mitral valve plane position is generated based on a low level edge detection using dynamic programming. It is shown that this method performs well in FRU and commercial matrix transducer data with low computational costs, which promises great value as an initialization step for segmentation approaches.

*automatic
initialization*

Chapter 6 aims at fulfilling the ultimate goal of our research objective, namely fully automatic segmentation of the left ventricle in 3D echocardiography. Our approach employs 3D active appearance models: statistical models that encompass typical shape and texture variations of the left ventricle in 3D echocardiography. The generalization of the shape and texture models is evaluated in relation to the number of patients in the study. For the texture model a nonlinear gray value normalization technique is presented, which improves the generalization capabilities of the texture model. Furthermore, the performance of the AAM detection is investigated with respect to the model's initialization, the a priori knowledge about the current patient and the optimization method used. It is shown that adequate automatic detection can be performed provided that the AAM is initialized within a certain lock-in range. Optimization using Jacobian tuning has shown to improve the matching accuracy and precision dramatically. In combination with automated initialization and proper postprocessing, 3D AAMs should be able to provide the desired automated analysis. This is a promising direction for future work.

*fully automatic
segmentation*

Finally, chapter 7 briefly presents the research objective with the conclusions

discussion and conclusions of each of the chapters in this thesis. Furthermore, a general discussion on the presented work and recommendations for future work are given with respect to texture analysis, model-based segmentation and integration of the presented techniques. In the general conclusions the main contributions of this thesis are described, showing the relevance of this work, in reducing analysis time and interobserver variability, and improving the reducibility by automating the analysis of 3D echocardiography. Methods presented in this thesis may be of great value in related research areas such as freehand 3D ultrasound and 3D transesophageal echocardiography. Also research on automated analysis of cardiac MR or CT may benefit from the presented advances in analysis of 3D echocardiography, as 3D echocardiography does also benefit from advances on those fields. Integration of the presented techniques and technical improvements in 3D ultrasound imaging may soon lead to reliable fully automated analysis of 3D echocardiography.

Samenvatting

9

Het linker ventrikel van het hart is verantwoordelijk voor het rondpompen van het zuurstofrijke bloed door het hele lichaam. Bepaling van de linkerventrikelfunctie is daarom erg belangrijk in het maken van keuzes in de kliniek, de evaluatie van therapeutische effecten en het bepalen van prognoses. Gedurende het laatste decennium jaar heeft 3-dimensionale (3D) echocardiografie zich ontpopt tot een zeer waardevolle techniek voor het evalueren van de linkerventrikelfunctie. Dit proefschrift, zoals gesteld in hoofdstuk 1, draagt bij aan het automatiseren van de analyse van 3D echocardiografie om de tijdsduur van de analyse te verkorten, variabiliteit tussen beoordelaars te verminderen en om de reproduceerbaarheid van de functiebepaling te vergroten. Op deze manier zouden de diagnose en de behandeling van cardiovasculaire aandoeningen verbeterd kunnen worden en de bijbehorende kosten worden gereduceerd. In dit hoofdstuk zal een overzicht van de beginselen van echografie gegeven worden. De huidige 3D echocardiografie, inclusief de fast rotating ultrasound (FRU) transducer en de commerciële matrix transducers, worden in een historisch perspectief geplaatst. Ook worden de belangrijkste onderwerpen in digitale beeldanalyse behandeld, met de nadruk op automatische analyse van echocardiografie.

Hoofdstuk 2 beschrijft een methode voor detectie van de endocardiale wand voor betrouwbare bepaling van het linkerventrikelvolumen over de volledige hartslag, terwijl slechts minimale gebruikersinteractie vereist is. In dit hoofdstuk ligt de nadruk op de analyse van beelden verkregen met de FRU transducer. De methode is gebaseerd op het detecteren van het endocardiale wand in 2-dimensionale (2D) doorsneden van het linker ventrikel, met behulp van *dynamic programming* en *pattern matching*. Vier handgetekende contouren in 2D doorsneden werden gebruikt om 3D plus tijd (3D+T) patiënt-specifieke vorm- en textuurmodellen af te leiden, die een benadering van de vorm en textuur van de endocardiale wand beschrijven. De kennis in deze modellen werd vervolgens optimaal benut met behulp van *dynamic programming* en *pattern matching* om de contouren in alle benodigde doorsneden in de gehele hartslag te detecteren. Met deze gedetecteerde contouren konden de 3D endocardial oppervlakken gereconstrueerd worden en het volume van het linker ventrikel over de volledige hartslag worden bepaald. Deze semi-automatische methode is geëvalueerd op data van 10 patiënten en vergeleken met volumes verkregen uit MRI opnamen, de gouden standaard. Het bleek dat de volumes goed correspondeerden met de MRI-bepaalde volumes. Een sterk punt van deze methode, de mogelijkheid om de automatisch gedetecteerde contouren te corrigeren, werd ook geëvalueerd, en liet een betere correlatie zien met de MRI volumes in vergelijking tot analyse zonder correcties. Ook, kwam er een lage variabiliteit uit de vergelijking van de resultaten van twee verschillende gebruikers.

Hoofdstuk 3 behandelt een belangrijk nadeel van de semi-automatische endocardiale wanddetectie: het beperkte model voor de beweging van de mitralisklep.

Daarom werd een techniek onderzocht voor het automatisch volgen van de mitralisklep schanierpunten (MKSP). Om een continue, cyclische detectie te garanderen, terwijl beweging in zowel horizontale als verticale richting werd toegestaan, werd het volgen van de MKSP gedaan met behulp van multidimensionaal *dynamic programming* met *apodized block matching*. De methode werd geëvalueerd op gedigitaliseerde 2D echocardiogrammen van doorsnee kwaliteit, om de robustheid van de tracing te beoordelen en om hem te vergelijken met gangbare detectietechnieken. Een hoge accuraatheid en precisie werd gevonden, vergelijkbaar met de variabiliteit tussen verschillende handmatige annotaties. Hierin is de grote waarde van deze methode te zien, en zijn belofte voor toepassing in 3D echocardiografie.

volgen van de
mitralisklep

In hoofdstuk 4 wordt de reconstructie van 3D+T data, verkregen met de FRU transducer, bestudeerd. De typische uitdagingen van dit reconstructieproces worden besproken. De gepresenteerde methode gebaseerd op *normalized convolution* is toegespitst op de onregelmatige en schaarse verdeling van de data. De aanpak werd gedemonstreerd op gesimuleerde 2D data en getest op in-vivo data verkregen met de FRU transducer. The voorgestelde reconstructietechniek laat voordelen zien ten opzichte van standaard technieken. Beeldartefacten als gevolg van fusie van opeenvolgende hartslagen en valse randen in dunbemonsterde delen van het beeld werden verminderd. De methode is generiek en kan worden toegepast op een breed scala van 3D ultrageluidsbeelden, zoals handmatige 3D echo en mechanische 3D transesophagale echocardiografie.

3D+T beeld-
reconstructie

Hoofdstuk 5 behandelt een oplossing voor een veelvoorkomend probleem in modelgebaseerde segmentatietechnieken: de initialisatie van het model. Daartoe wordt een detectietechniek gepresenteerd die een skeletachtige benadering van het linker ventrikel detecteert, namelijk de lange as van het linker ventrikel en de positie van het vlak door de mitralisklep. Deze twee elementen kunnen zeer behulpzaam zijn bij het initialiseren van een modelgebaseerde techniek. De gepresenteerde techniek kan de lange as betrouwbaar vinden met behulp van *dynamic programming* en een *Hough transform* voor cirkels, gebaseerd op detectie van het midden van het linker ventrikel in vlakken die loodrecht op de geschatte lange as staan en door deze middelpunten te volgen over de volledige hartslag. Door gebruik te maken van deze gedetecteerde lange as kan een benadering gevonden worden voor het vlak door de mitralisklep, gebaseerd op basale contourdetectie met *dynamic programming*. Er wordt gedemonstreerd dat de gepresenteerde techniek goed presteert op FRU beelddata en ook in beelden verkregen met een commerciële matrix-transducer. De techniek vraagt weinig rekentijd en is daarom veelbelovend om als initialisatiestap voor segmentatiemethoden gebruikt te worden.

automatische
initialisatie

Hoofdstuk 6 richt zich op het vervullen van het uiteindelijke doel binnen dit onderzoek, namelijk volautomatische segmentatie van het linker ventrikel in 3D echocardiografie. We kozen voor een aanpak met gebaseerd op 3D *active appear-*

volautomatische
segmentatie

ance models (AAMs): statistische modellen die typische variatie in vorm en textuur van het linker ventrikel beschrijven in 3D echocardiografie. Het generaliserend vermogen van de vorm- en textuurmodellen werd geëvalueerd in relatie tot het aantal patienten in de studie. Voor het textuurmodel wordt een niet-lineaire grijswaarde-normalisatie gepresenteerd, die de generalisatie van het textuurmodel verbetert. Verder wordt de prestatie van de AAM-detectie uitgebreid geëvalueerd in relatie tot de initialisatie van het model, de hoeveelheid vooraf bekende informatie over het te segmenteren beeld en de gebruikte optimalisatiemethode. Dit laat zien dat een adequate automatische detectie kan worden bereikt mits de AAM is geïnitieerd binnen een bepaald insluitingsgebied. Het gebruik van *Jacobian tuning* in de optimalisatie blijkt een aanzienlijke verbetering te geven in de precisie en accuratesse van de detectie. Wanneer 3D AAMs gecombineerd gebruikt worden met een automatische initialisatie en een passende nabewerking, dan zouden ze in staat moeten zijn om de gewenste analyse te leveren. Dit is een veelbelovende richting voor toekomstig onderzoek.

discussie en
conclusies

Tenslotte presenteert hoofdstuk 7 in het kort het onderzoeksdoel met de conclusies van elk van de hoofdstukken in dit proefschrift. Ook wordt het gepresenteerde werk bediscussieerd en worden aanbevelingen voor toekomstig onderzoek gegeven, in het bijzonder met het oog op textuuranalyse, modelgebaseerde segmentatie en integratie van de gepresenteerde technieken. De belangrijkste bijdragen en de relevantie van het werk worden samengevat in de algemene conclusies. Daar komen de reductie in analysetijd en variabiliteit tussen gebruikers aan bod, en het verbeteren van de reproduceerbaarheid door het automatiseren van de analyse van 3D echocardiografie. De gepresenteerde technieken kunnen ook van grote waarde zijn in gerelateerde toepassingsgebieden, zoals handmatige 3D echografie en 3D transesophagale echocardiografie. Ook onderzoek op het gebied van automatische analyse van cradiale MR of CT zou kunnen profiteren van de voorgestelde technieken, net zoals onderzoek op het gebied van 3D echocardiografie profiteert van ontwikkeling in die onderzoeksrichtingen. Integratie van de gepresenteerde technieken en technische innovaties op het gebied van 3D echografie kunnen snel leiden tot betrouwbare volautomatische analyse van 3D echocardiografie.

Bibliography

- Achim, A., A. Bezerianos, and P. Tsakalides [2001]. Novel Bayesian multiscale method for speckle removal in medical ultrasound images. *IEEE T Med Imaging* 20;8: 772–83.
- Adam, D., S. Beilin-Nissan, Z. Friedman, and V. Behar [2006]. The combined effect of spatial compounding and nonlinear filtering on the speckle reduction in ultrasound images. *Ultrasonics* 44;2: 166–81.
- Alam, M. and C. Hoglund [1992]. Assessment by echocardiogram of left ventricular diastolic function in healthy subjects using the atrioventricular plane displacement. *Am J Cardiol* 69;5: 565–8.
- Amini, A., T. E. Weymouth, and R. C. Jain [1990]. Using dynamic programming for solving variational problems in vision. *IEEE T Pattern Anal* 12;9: 855–867.
- Angelini, E. D., A. F. Laine, S. Takuma, J. W. Holmes, and S. Homma [2001]. LV volume quantification via spatiotemporal analysis of real-time 3-D echocardiography. *IEEE T Med Imaging* 20;6: 457–69.
- Angelini, E. D., S. Homma, G. Pearson, J. W. Holmes, and A. F. Laine [2005]. Segmentation of real-time three-dimensional ultrasound for quantification of ventricular function: a clinical study on right and left ventricles. *Ultrasound Med Biol* 31;9: 1143–1158.
- Armstrong, W. F. and W. A. Zoghbi [2005]. Stress echocardiography: current methodology and clinical applications. *J Am Coll Cardiol* 45;11: 1739–47.
- Ballard, D. H. [1981]. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recogn* 13;2: 111–122.
- Barnett, S. B., G. R. Ter Haar, M. C. Ziskin, H. D. Rott, F. A. Duck, and K. Maeda [2000]. International recommendations and guidelines for the safe use of diagnostic ultrasound in medicine. *Ultrasound Med Biol* 26;3: 355–66.

- Barratt, D. C., A. H. Davies, A. D. Hughes, S. A. Thom, and K. N. Humphries [2001]. Optimisation and evaluation of an electromagnetic tracking device for high-accuracy three-dimensional ultrasound imaging of the carotid arteries. *Ultrasound Med Biol* 27;7: 957–968.
- Batur, A. U. and M. H. Hayes [2005]. Adaptive active appearance models. *IEEE T Image Process* 14;11: 1707–21.
- Bax, J. J., J. A. Patton, D. Poldermans, A. Elhendy, and M. P. Sandler [2000]. 18-Fluorodeoxyglucose imaging with positron emission tomography and single photon emission computed tomography: cardiac applications. *J Nucl Med* 30;4: 281–98.
- Behar, V., D. Adam, P. Lysyansky, and Z. Friedman [2004]. The combined effect of nonlinear filtration and window size on the accuracy of tissue displacement estimation using detected echo signals. *Ultrasonics* 41;9: 743–53.
- Beichel, R., G. Gotschuli, E. Sorantin, F. Leberl, and M. Sonka [2002]. Diaphragm dome surface segmentation in CT data sets: a 3D active appearance model approach. *Proc SPIE Med Imaging* 4684: 475–484.
- Beichel, R., H. Bischof, F. Leberl, and M. Sonka [2005]. Robust active appearance models and their application to medical image analysis. *IEEE T Med Imaging* 24;9: 1151–1169.
- Bellmann, R. E. [1965]. *Dynamic programming*. Princeton Univ Pr.
- Belohlavek, M., K. Tanabe, D. Jakrapanichakul, J. F. Breen, and J. B. Seward [2001]. Rapid three-dimensional echocardiography: clinically feasible alternative for precise and accurate measurement of left ventricular volumes. *Circulation* 103;24: 2882–4.
- Binder, T., M. Sussner, D. Moertl, T. Strohmer, H. Baumgartner, G. Maurer, and G. Porenta [1999]. Artificial neural networks and spatial temporal contour linking for automated endocardial contour detection on echocardiograms: a novel approach to determine left ventricular contractile function. *Ultrasound Med Biol* 25;7: 1069–76.
- Bland, J. M. and D. G. Altman [1986]. Statistical methods for assessing agreement between two methods of clinical measurement. *Lancet* 1;8476: 307–10.
- Bohs, L. N. and G. E. Trahey [1991]. A novel method for angle independent ultrasonic imaging of blood flow and tissue motion. *IEEE T Bio-med Eng* 38;3: 280–6.
- Bosch, H. G., G. van Burken, F. Nijland, and J. H. C. Reiber [1998]. Overview of automated quantitation techniques in 2D echocardiography. In: *What's New in Cardiovascular Imaging*. Ed. by J. H. C. Reiber and E. E. van der Wall. Springer.
- Bosch, J. G. [2007]. Echocardiographic digital image processing and approaches to automated border detection. In: *The practice of clinical echocardiography*. Ed. by C. M. Otto. 3rd ed. Philadelphia: Saunders. 262–282.

- Bosch, J. G., S. C. Mitchell, B. P. F. Lelieveldt, M. Sonka, F. Nijland, and J. H. C. Reiber [2000]. Model-based automated border detection for quantitative stress echocardiography. *Eur Heart J* 21;Abstract suppl; 37.
- Bosch, J. G., S. C. Mitchell, B. P. F. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J. H. C. Reiber [2002]. Automatic segmentation of echocardiographic sequences by active appearance motion models. *IEEE T Med Imaging* 21;11; 1374–1383.
- Broyden, C. G. [1965]. A class of methods for solving nonlinear simultaneous equations. *Math Comput* 19;92; 577–593.
- Bruining, N., C. T. Lancée, J. R. T. C. Roelandt, and N. Bom [2000]. Three-dimensional echocardiography paves the way toward virtual reality. *Ultrasound Med Biol* 26;7; 1065–1074.
- Canals, R., G. Lamarque, and P. Chatain [1999]. Volumetric ultrasound system for left ventricle motion imaging. *IEEE T Ultrason Ferr* 46;6; 1527–1538.
- Chen, E. J., W. K. Jenkins, and W. D. O'Brien Jr [1995]. Performance of ultrasonic speckle tracking in various tissues. *J Acoust Soc Am* 98;3; 1273–8.
- Christodoulou, C. I., C. S. Pattichis, M. Pantziaris, and A. Nicolaides [2003]. Texture-based classification of atherosclerotic carotid plaques. *IEEE T Med Imaging* 22;7; 902–12.
- Comaniciu, D., X. S. Zhou, and S. Krishnan [2004]. Robust real-time myocardial border tracking for echocardiography: an information fusion approach. *IEEE T Med Imaging* 23;7; 849–60.
- Cootes, T. F. and P. Kittipanya-ngam [2002]. Comparing variations on the active appearance model algorithm. *Proc Brit Mach Vis Conf*. 837–846.
- Cootes, T. F. and C. J. Taylor [2001a]. *Statistical models of appearance for computer vision*. http://www.isbe.man.ac.uk/~bim/Models/app_model.ps.gz.
- [2001b]. Statistical models of appearance for medical image analysis and computer vision. *Proc SPIE Med Imaging* 4322; 236–248.
- [2006]. An algorithm for tuning an active appearance model to new data. *Proc Brit Mach Vis Conf* 3; 919–928.
- Corbett, J. R. and E. P. Ficaro [1999]. Clinical review of attenuation-corrected cardiac SPECT. *J Nucl Cardiol* 6;1 Pt 1; 54–68.
- Corsi, C., G. Saracino, A. Sarti, and C. Lamberti [2002]. Left ventricular volume estimation for real-time three-dimensional echocardiography. *IEEE T Med Imaging* 21;9; 1202–1208.
- DeCara, J. M., E. Toledo, I. S. Salgo, G. Lammertin, L. Weinert, and R. M. Lang [2005]. Evaluation of left ventricular systolic function using automated angle-independent motion tracking of mitral annular displacement. *J Am Soc Echocardiogr* 18;12; 1266–9.
- Dekker, D. L., R. L. Piziali, and E. Dong Jr [1974]. A system for ultrasonically imaging the human heart in three dimensions. *Comput Biomed Res* 7;6; 544–53.
-

- Delabays, A., N. G. Pandian, Q. L. Cao, L. Sugeng, G. Marx, A. Ludomirski, and S. L. Schwartz [1995]. Transthoracic real-time three-dimensional echocardiography using a fan-like scanning approach for data acquisition: methods, strengths, problems, and initial clinical experience. *Echocardiography* 12;1; 49–59.
- Djoa, K. K., N. de Jong, F. C. van Egmond, J. D. Kasprzak, W. B. Vletter, C. T. Lancee, A. F. W. van der Steen, N. Bom, and J. R. T. C. Roelandt [2000]. A fast rotating scanning unit for real-time three-dimensional echo data acquisition. *Ultrasound Med Biol* 26;5; 863–9.
- Drori, I., D. Cohen-Or, and H. Yeshurun [2003]. Fragment-based image completion. *ACM T Graphic* 22;3; 303–312.
- Duck, F. A. [2002]. Nonlinear acoustics in diagnostic ultrasound. *Ultrasound Med Biol* 28;1; 1–18.
- Engelse, W. A. H. and C. Zeelenberg [1979]. A single scan algorithm for QRS-detection and feature extraction. *Proc Comput Card* 6; 37–42.
- Estépar, R. S. J., M. Martín-Fernández, C. Alberola-López, J. Ellsmere, R. Kikinis, and C. E. Westin [2003]. Freehand ultrasound reconstruction based on ROI prior modeling and normalized convolution. *Lect Notes Comput Sc* 2879; 382–390.
- Eto, Y., H. Yamada, J. H. Shin, D. A. Agler, H. Tsujino, J. X. Qin, G. Saracino, N. L. Greenberg, J. D. Thomas, and T. Shiota [2005]. Automated mitral annular tracking: a novel method for evaluating mitral annular motion using two-dimensional echocardiography. *J Am Soc Echocardiog* 18;4; 306–12.
- Farsiu, S., D. Robinson, M. Elad, and P. Milanfar [2004]. Advances and challenges in super-resolution. *Int J Imag Syst Tech* 14;2; 47–57.
- Feigenbaum, H. [2004]. *Feigenbaum's echocardiography*. 6th ed. Philadelphia, PA: Lippincott Williams & Wilkins.
- Flachskampf, F. A., S. Chandra, A. Gaddipatti, R. A. Levine, A. E. Weyman, W. Ameling, P. Hanrath, and J. D. Thomas [2000]. Analysis of shape and motion of the mitral annulus in subjects with and without cardiomyopathy by echocardiographic 3-dimensional reconstruction. *J Am Soc Echocardiog* 13;4; 277–87.
- Frangi, A. F., W. J. Niessen, and M. A. Viergever [2001]. Three-dimensional modeling for functional analysis of cardiac images, a review. *IEEE T Med Imaging* 20;1; 2–25.
- Friedland, N. and D. Adam [1989]. Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing. *IEEE T Med Imaging* 8;4; 344–53.
- Funda, J. and R. P. Paul [1988]. A comparison of transforms and quaternions in robotics. *Proc IEEE Int Conf Robo Automat* 2; 886–891.
- Geleijnse, M. L., P. M. Fioretti, and J. R. T. C. Roelandt [1997]. Methodology, feasibility, safety and diagnostic accuracy of dobutamine stress echocardiography. *J Am Coll Cardiol* 30;3; 595–606.

- Gérard, O., A. C. Billon, J.-M. Rouet, M. Jacob, M. Fradkin, and C. Allouche [2002]. Efficient model-based quantification of left ventricular function in 3-D echocardiography. *IEEE T Med Imaging* 21;9: 1059–1068.
- Golemati, S., A. Sassano, M. J. Lever, A. A. Bharath, S. Dhanjil, and A. N. Nicolaides [2003]. Carotid artery wall motion estimated from B-mode ultrasound using region tracking and block matching. *Ultrasound Med Biol* 29;3: 387–99.
- Goodall, C. [1991]. Procrustes methods in the statistical analysis of shape. *J Roy Stat Soc B Met* 53;2: 285–339.
- Gower, J. C. [1975]. Generalized procrustes analysis. *Psychometrika* 40;1: 33–51.
- Gross, R., I. Matthews, and S. Baker [2006]. Active appearance models with occlusion. *Image Vision Comput* 24;6: 593–604.
- Hansegård, J., S. Urheim, K. Lunde, and S. I. Rabben [2007a]. Constrained active appearance models for segmentation of triplane echocardiograms. *IEEE T Med Imaging* 26;10: 1391–400.
- Hansegård, J., F. Orderud, and S. I. Rabben [2007b]. Real-time active shape models for segmentation of 3D cardiac ultrasound. *Lect Notes Comput Sc* 4673: 157–164.
- Horn, B. K. P. [1987]. Closed-form solution of absolute orientation using unit quaternions. *J Opt Soc Am* 4;4: 629–642.
- Jacob, G., J. A. Noble, C. Behrenbruch, A. D. Kelion, and A. P. Banning [2002]. A shape-space-based approach to tracking myocardial borders and quantifying regional left-ventricular function applied in echocardiography. *IEEE T Med Imaging* 21;3: 226–38.
- Jenkins, C., K. Bricknell, L. Hanekom, and T. H. Marwick [2004]. Reproducibility and accuracy of echocardiographic measurements of left ventricular parameters using real-time three-dimensional echocardiography. *J Am Coll Cardiol* 44;4: 878–86.
- Kaus, M. R., J. von Berg, J. Weese, W. J. Niessen, and V. Pekar [2004]. Automated segmentation of the left ventricle in cardiac MRI. *Med Image Anal* 8;3: 245–254.
- Knutsson, H. and C. F. Westin [1993]. Normalized and differential convolution. *Proc IEEE Int Conf Comput Vis Pattern Recogn.* 515–523.
- Kühl, H. P., M. Schreckenberg, D. Rulands, M. Katoh, W. Schafer, G. Schummers, A. Buckner, P. Hanrath, and A. Franke [2004]. High-resolution transthoracic real-time three-dimensional echocardiography: quantitation of cardiac volumes and function using semi-automatic border detection and comparison with cardiac magnetic resonance imaging. *J Am Coll Cardiol* 43;11: 2083–90.
- Larsen, R., M. B. Stegmann, S. Darkner, S. Forchhammer, T. F. Cootes, and B. Kjær Ersbøll [2007]. Texture enhanced appearance models. *Comput Vis Image Und* 106;1: 20–30.
-

- Leung, K. Y. E., R. A. Baldewsing, F. Mastik, J. A. Schaar, A. Gisolf, and A. F. W. Van der Steen [2006a]. Motion compensation for intravascular ultrasound palpography. *IEEE T Ultrason Ferr* 53;7; 1269–1280.
- Leung, K. Y. E., M. van Stralen, G. van Burken, M. M. Voormolen, A. Nemes, F. J. ten Cate, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2006b]. Sparse appearance model based registration of 3D ultrasound images. *Lect Notes Comput Sc* 4091; 236–243.
- Loke, R. E. and J. M. H. du Buf [2004]. Quadtree-guided 3-D interpolation of irregular sonar data sets. *IEEE J Oceanic Eng* 29;2; 457–471.
- Lorenz, C. H., E. S. Walker, V. L. Morgan, S. S. Klein, and Jr. Graham T. P. [1999]. Normal human right and left ventricular mass, systolic function, and gender differences by cine magnetic resonance imaging. *J Cardio Magn Reson* 1;1; 7–21.
- Mannaerts, H. E., J. A. van der Heide, O. Kamp, T. Papavassiliu, J. T. Marcus, A. Beek, A. C. van Rossum, J. Twisk, and C. A. Visser [2003]. Quantification of left ventricular volumes and ejection fraction using freehand transthoracic three-dimensional echocardiography: comparison with magnetic resonance imaging. *J Am Soc Echocardiogr* 16;2; 101–9.
- Meairs, S., J. Beyer, and M. Hennerici [2000]. Reconstruction and visualization of irregularly sampled three- and four-dimensional ultrasound data for cerebrovascular applications. *Ultrasound Med Biol* 26;2; 263–72.
- Mignotte, M. and J. Meunier [2001]. A multiscale optimization approach for the dynamic contour-based boundary detection issue. *Comput Med Imag Grap* 25;3; 265–75.
- Mitchell, S. C., B. P. F. Lelieveldt, R. J. van der Geest, J. G. Bosch, J. H. C. Reiber, and M. Sonka [2001a]. Multistage hybrid active appearance model matching: segmentation of left and right ventricles in cardiac MR images. *IEEE T Med Imaging* 20;5; 415–423.
- [2001b]. Time continuous segmentation of cardiac MR image sequences using active appearance motion models. *Proc SPIE Med Imaging* 4322; 249–256.
- Mitchell, S. C., J. G. Bosch, B. P. F. Lelieveldt, R. J. van der Geest, J. H. C. Reiber, and M. Sonka [2002]. 3-D active appearance models: segmentation of cardiac MR and ultrasound images. *IEEE T Med Imaging* 21;9; 1167–1178.
- Mittrapiyanuruk, P., G. N. DeSouza, and A. C. Kak [2005]. Accurate 3D tracking of rigid objects with occlusion using active appearance models. *Proc IEEE Workshop Motion Vis Comput*.
- Montagnat, J. and H. Delingette [2000]. Space and time shape constrained deformable surfaces for 4D medical image segmentation. *Lect Notes Comput Sc* 1935; 687–696.
- Moritz, W. E. and P. L. Shreve [1976]. A microprocessor-based spatial-locating system for use with diagnostic ultrasound. *P IEEE* 64;6; 966–974.

- Müller, A., A. Neitmann, N. Merkle, J. Wohrle, V. Hombach, and H. A. Kestler [2005]. Contour detection of short axis slice MR images for contraction irregularity assessment. *Proc Comput Card* 32; 21–24.
- Nijland, E., O. Kamp, P. M. Verhorst, W. G. de Voogt, H. G. Bosch, and C. A. Visser [2002]. Myocardial viability: impact on left ventricular dilatation after acute myocardial infarction. *Heart* 87;1; 17–22.
- Nillesen, M. M., R. G. Lopata, I. H. Gerrits, L. Kapusta, H. J. Huisman, J. M. Thijssen, and C. L. de Korte [2007]. Segmentation of the heart muscle in 3-D pediatric echocardiographic images. *Ultrasound Med Biol* 33;9; 1453–62.
- Nillesen, M. M., R. G. P. Lopata, I. H. Gerrits, L. Kapusta, J. M. Thijssen, and C. L. de Korte [2008]. Modeling envelope statistics of blood and myocardium for segmentation of echocardiographic images. *Ultrasound Med Biol* 34;4; 674–680.
- Noble, J. A. and D. Boukerroui [2006]. Ultrasound image segmentation: a survey. *IEEE T Med Imaging* 25;8; 987–1010.
- Nosir, Y. F., P. M. Fioretti, W. B. Vletter, E. Boersma, A. Salustri, J. T. Postma, A. E. Reijs, F. J. Ten Cate, and J. R. Roelandt [1996]. Accurate measurement of left ventricular ejection fraction by three-dimensional echocardiography. A comparison with radionuclide angiography. *Circulation* 94;3; 460–6.
- Nosir, Y. F., J. Stoker, J. D. Kasprzak, M. H. Lequin, A. Dall’Agata, F. J. Ten Cate, and J. R. Roelandt [1999]. Paraplane analysis from precordial three-dimensional echocardiographic data sets for rapid and accurate quantification of left ventricular volume and function: a comparison with magnetic resonance imaging. *Am Heart J* 137;1; 134–43.
- Oost, E., G. Koning, M. Sonka, P. V. Oemrawsingh, J. H. C. Reiber, and B. P. F. Lelieveldt [2006]. Automated contour detection in X-ray left ventricular angiograms using multiview active appearance models and dynamic programming. *IEEE T Med Imaging* 25;9; 1158–1171.
- Orderud, E., J. Hansgård, and S. I. Rabben [2007]. Real-time tracking of the left ventricle in 3D echocardiography using a state estimation approach. *Lect Notes Comput Sc* 4791; 858–865.
- Pai, R. G., M. M. Bodenheimer, S. M. Pai, J. H. Koss, and R. D. Adamick [1991]. Usefulness of systolic excursion of the mitral annulus as an index of left ventricular systolic function. *Am J Cardiol* 67;2; 222–4.
- Pandian, N. G., J. R. T. C. Roelandt, N. C. Nanda, L. Sugeng, Q. L. Cao, J. Azevedo, S. L. Schwartz, M. A. Vannan, A. Ludomirski, G. Marx, and M. Vogel [1994]. Dynamic three-dimensional echocardiography: methods and clinical potential. *Echocardiography* 11;3; 237–259.
- Papavassiliou, D., N. R. Doelling, M. K. Bowman, H. Yeung, J. Rock, B. Klas, K. Chung, and D. A. Fyfe [1998]. Initial experience with an internally rotating transthoracic
-

- three-dimensional echocardiographic probe and image acquisition on a conventional echocardiogram machine. *Echocardiography* 15;4: 369–376.
- Penney, G. P., J. A. Schnabel, D. Rueckert, M. A. Viergever, and W. J. Niessen [2004]. Registration-based interpolation. *IEEE T Med Imaging* 23;7: 922–6.
- Pham, T. Q. and L. J. van Vliet [2003]. Normalized averaging using adaptive applicability functions with applications in image reconstruction from sparsely and randomly sampled data. *Proc Scand Conf Image Anal.* 485–492.
- Raab, F. H., E. B. Blood, T. O. Steiner, and H. R. Jones [1979]. Magnetic position and orientation tracking system. *IEEE T Aero Elec Sys* AES-15;5: 709–718.
- Rabben, S. I., A. H. Torp, A. Støylen, S. Slørdahl, K. Bjørnstad, B. O. Haugen, and B. Angelsen [2000]. Semiautomatic contour detection in ultrasound M-mode images. *Ultrasound Med Biol* 26;2: 287–96.
- Roberts, M., T. F. Cootes, and J. E. Adams [2003]. Linking sequences of active appearance sub-models via constraints: an application in automated vertebral morphometry. *Proc Brit Mach Vis Conf* 1; 349–358.
- Roberts, M., T. F. Cootes, E. Pacheco, and J. E. Adams [2007]. Quantitative vertebral fracture detection on DXA images using shape and appearance models. *Acad Radiol* 14;10: 1166–1178.
- Sanchez-Ortiz, G. I., G. J. T. Wright, N. Clarke, J. Declerck, A. P. Banning, and J. A. Noble [2002]. Automated 3-D echocardiography analysis compared with manual delineations and SPECT MUGA. *IEEE T Med Imaging* 21;9: 1069–1076.
- Schuijff, J. D. and J. J. Bax [2008]. CT angiography: an alternative to nuclear perfusion imaging? *Heart* 94;3: 255–7.
- Schwartz, S. L., Q. L. Cao, J. Azevedo, and N. G. Pandian [1994]. Simulation of intraoperative visualization of cardiac structures and study of dynamic surgical anatomy with real-time three-dimensional echocardiography. *Am J Cardiol* 73;7: 501–507.
- Sivaramakrishna, R., K. A. Powell, M. L. Lieber, W. A. Chilcote, and R. Shekhar [2002]. Texture analysis of lesions in breast ultrasound images. *Comput Med Imag Grap* 26;5: 303–7.
- Smith, S. W., H. G. Pavy Jr, and O. T. von Ramm [1991]. High-speed ultrasound volumetric imaging system. I. Transducer design and beam steering. *IEEE T Ultrason Ferr* 38;2: 100–108.
- Sonka, M., V. Hlavac, and R. Boyle [1999]. *Image processing, analysis, and machine vision*. 2nd ed. Pacific Grove, CA: Brooks/Cole Publishing company.
- Spencer, K. T., J. Bednarz, P. G. Rafter, C. Korcarz, and R. M. Lang [1998]. Use of harmonic imaging without echocardiographic contrast to improve two-dimensional image quality. *Am J Cardiol* 82;6: 794–799.

- Stegmann, M. B. and D. Pedersen [2005]. Bi-temporal 3 D active appearance models with applications to unsupervised ejection fraction estimation. *Proc SPIE Med Imaging* 5747; 336–350.
- Stegmann, M. B., B. K. Ersboll, and R. Larsen [2003]. FAME—a flexible appearance modeling environment. *IEEE T Med Imaging* 22;10; 1319–31.
- Stegmann, M. B., H. Olafsdottir, and H. B. W. Larsson [2005]. Unsupervised motion-compensation of multi-slice cardiac perfusion MRI. *Med Image Anal* 9;4; 394–410.
- Stetten, G. D. and S. M. Pizer [1999]. Medial-node models to identify and measure objects in real-time 3-D echocardiography. *IEEE T Med Imaging* 18;10; 1025–1034.
- Tauber, C., H. Batatia, and A. Ayache [2004]. A robust speckle reducing anisotropic diffusion. *Proc Int Conf Image Process* 1; 247–250.
- Thodberg, H. H. [2002]. Hands-on experience with active appearance models. *Proc SPIE Med Imaging* 4684; 495–506.
- Thomas, J. D. and D. N. Rubin [1998]. Tissue harmonic imaging: why does it work? *J Am Soc Echocardiog* 11;8; 803–808.
- Tranquart, F., N. Grenier, V. Eder, and L. Pourcelot [1999]. Clinical use of ultrasound tissue harmonic imaging. *Ultrasound Med Biol* 25;6; 889–894.
- Unser, M., G. Pelle, P. Brun, and M. Eden [1989]. Automated extraction of serial myocardial borders from M-mode echocardiograms. *IEEE T Med Imaging* 8;1; 96–103.
- Üzümcü, M., R. J. van der Geest, M. Sonka, H. J. Lamb, J. H. C. Reiber, and B. P. F. Lelieveldt [2005]. Multiview active appearance models for simultaneous segmentation of cardiac 2- and 4-chamber long-axis magnetic resonance images. *Invest Radiol* 40;4; 195–203.
- Üzümcü, M., R. J. van der Geest, C. Swingen, J. H. C. Reiber, and B. P. F. Lelieveldt [2006]. Time continuous tracking and segmentation of cardiovascular magnetic resonance images using multidimensional dynamic programming. *Invest Radiol* 41;1; 52–62.
- Van der Geest, R. J., V. G. Buller, E. Jansen, H. J. Lamb, L. H. Baur, E. E. van der Wall, A. de Roos, and J. H. Reiber [1997]. Comparison between manual and semiautomated analysis of left ventricular volume parameters from short-axis MR images. *J Comput Assist Tomo* 21;5; 756–65.
- Van Stralen, M., J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, R. M. van Geuns, C. T. Lancée, N. de Jong, and J. H. C. Reiber [2005]. Left ventricular volume estimation in cardiac three-dimensional ultrasound: a semiautomatic border detection approach. *Acad Radiol* 12;10; 1241–1249.
- Varandas, J., P. Baptista, J. Santos, R. Martins, and J. Dias [2004]. VOLUS—a visualization system for 3D ultrasound data. *Ultrasonics* 42;1-9; 689–694.
-

- Veronesi, F., C. Corsi, E. G. Caiani, A. Sarti, and C. Lamberti [2006]. Tracking of left ventricular long axis from real-time three-dimensional echocardiography using optical flow techniques. *IEEE T Inf Technol B* 10;1: 174–81.
- Von Ramm, O. T., S. W. Smith, and H. G. Pavy Jr [1991]. High-speed ultrasound volumetric imaging system. II. Parallel processing and image display. *IEEE T Ultrason Ferr* 38;2: 109–115.
- Voormolen, M. M. and M. G. Danilouchkine [2007]. Aspects of left ventricular volume comparison between 3D echocardiography and MRI. *J Am Soc Echocardiog* 20;12: 1421–1422.
- Voormolen, M. M., A. Bouakaz, B. J. Krenning, C. T. Lancée, F. J. ten Cate, J. R. T. C. Roelandt, A. F. van der Steen, and N. de Jong [2002]. A new transducer for 3D harmonic imaging. *Proc IEEE Int Ultrason Symp.* 1261–1264.
- Voormolen, M. M., B. J. Krenning, C. T. Lancée, F. J. ten Cate, J. R. T. C. Roelandt, A. F. W. van der Steen, and N. de Jong [2003]. Quantitative harmonic 3D echocardiography with a fast rotating ultrasound transducer. *Proc IEEE Int Ultrason Symp.* 122–125.
- Voormolen, M. M., B. J. Krenning, C. T. Lancée, F. J. ten Cate, J. R. Roelandt, A. F. van der Steen, and N. de Jong [2006]. Harmonic 3-D echocardiography with a fast-rotating ultrasound transducer. *IEEE T Ultrason Ferr* 53;10: 1739–48.
- Voormolen, M. M., B. J. Krenning, R. J. van Geuns, J. Borsboom, C. T. Lancée, F. J. ten Cate, J. R. T. C. Roelandt, A. F. van der Steen, and N. de Jong [2007]. Efficient quantification of the left ventricular volume using 3-dimensional echocardiography: the minimal number of equiangular long-axis images for accurate quantification of the left ventricular volume. *J Am Soc Echocardiog* 20;4: 373–80.
- WHO, World Health Organisation [2007]. Fact sheet N° 317: Cardiovascular disease. <http://www.who.int/mediacentre/factsheets/fs317/en/index.html>.
- Ward, B., A. C. Baker, and V. F. Humphrey [1997]. Nonlinear propagation applied to the improvement of resolution in diagnostic medical ultrasound. *J Acoust Soc Am* 101;1: 143–154.
- Wells, P. N. T. and M. Halliwell [1981]. Speckle in ultrasonic imaging. *Ultrasonics* 19;5: 225–229.
- Westin, C. F. [1994]. A tensor framework for multidimensional signal processing. PhD thesis. Linköping University.
- Willenheimer, R., B. Israelsson, C. Cline, E. Rydberg, K. Broms, and L. Erhardt [1999]. Left atrioventricular plane displacement is related to both systolic and diastolic left ventricular performance in patients with chronic heart failure. *Eur Heart J* 20;8: 612–8.
- Williams, B. R. [1994]. A retrospective study of the diagnostic accuracy of a community hospital-based PET center for the detection of coronary artery disease using rubidium-82. *J Nucl Med* 35;10: 1586–1592.

- Xiao, C.-Y., S. Zhang, and Y.-Z. Chen [2004]. A diffusion stick method for speckle suppression in ultrasonic images. *Pattern Recogn Lett* 25;16; 1866–1876.
- Xie, J., Y. Jiang, and H. T. Tsui [2005]. Segmentation of kidney from ultrasound images based on texture and shape priors. *IEEE T Med Imaging* 24;1; 45–57.
- Xu, C. X. [1990]. Hybrid method for nonlinear least-square problems without calculating derivatives. *J Optimiz Theory App* 65;3; 555–574.
- Yoshida, H., D. D. Casalino, B. Keserci, A. Coskun, O. Ozturk, and A. Savranlar [2003]. Wavelet-packet-based texture analysis for differentiation between benign and malignant liver tumours in ultrasound images. *Phys Med Biol* 48;22; 3735–53.
- Yu, Y. and S. T. Acton [2002]. Speckle reducing anisotropic diffusion. *IEEE T Image Process* 11;11; 1260–1270.
- Zagrodsky, V., V. Walimbe, C. R. Castro-Pareja, J. X. Qin, J. M. Song, and R. Shekhar [2005]. Registration-assisted segmentation of real-time 3-D echocardiographic data using deformable models. *IEEE T Med Imaging* 24;9; 1089–99.
- Zheng, Y., A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu [2007]. Fast automatic heart chamber segmentation from 3D CT data using marginal space learning and steerable features. *Proc IEEE Int Conf Comput Vis.* 1–8.
- Zhou, X. S., D. Comaniciu, and A. Gupta [2005]. An information fusion framework for robust shape tracking. *IEEE T Pattern Anal* 27;1; 115–29.
- Zhu, Y., X. Papademetris, A. Sinusas, and J. S. Duncan [2007]. Segmentation of myocardial volumes from real-time 3D echocardiography using an incompressibility constraint? *Lect Notes Comput Sc* 4791; 44–51.
-

Publications

Journal papers

- Leung, K. Y. E., **M. van Stralen**, A. Nemes, M. M. Voormolen, G. van Burken, M. L. Geleijnse, F. J. Ten Cate, J. H. C. Reiber, N. de Jong, A. F. W. van der Steen, and J. G. Bosch [2008]. Sparse registration for three-dimensional stress echocardiography. *IEEE T Med Imaging* 27;11; 1568–79.
- Nemes, A., K. Y. E. Leung, G. van Burken, **M. van Stralen**, J. G. Bosch, O. I. Soliman, B. J. Krenning, W. B. Vletter, F. J. Cate, and M. L. Geleijnse [2008, in press]. Side-by-side viewing of anatomically aligned left ventricular segments in three-dimensional stress echocardiography. *Echocardiography*.
- Nevo, S. T., **M. van Stralen**, A. M. Vossepoel, J. H. Reiber, N. de Jong, A. F. W. van der Steen, and J. G. Bosch [2007]. Automated tracking of the mitral valve annulus motion in apical echocardiographic images using multidimensional dynamic programming. *Ultrasound Med Biol* 33;9; 1389–99.
- van Stralen, M.**, J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, R. M. van Geuns, C. T. Lancée, N. de Jong, and J. H. C. Reiber [2005]. Left ventricular volume estimation in cardiac three-dimensional ultrasound: a semiautomatic border detection approach. *Acad Radiol* 12;10; 1241–1249.
- van Stralen, M.**, K. Y. Leung, M. M. Voormolen, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2008]. Time continuous detection of the left ventricular long axis and the mitral valve plane in 3-D echocardiography. *Ultrasound Med Biol* 34;2; 196–207.

| Peer-reviewed conference proceeding papers

Leung, K. Y. E., **M. van Stralen**, G. van Burken, M. M. Voormolen, A. Nemes, F. J. ten Cate, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2006]. Sparse appearance model based registration of 3D ultrasound images. *Proc Med Imaging Augmented Reality Lect Notes Comput Sc* 4091; 236–243.

van Stralen, M., J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, C. T. Lancée, N. de Jong, and J. H. C. Reiber [2004]. A semi-automatic endocardial border detection method for the left ventricle in 4D ultrasound data sets. *Proc Medical Image Computing and Computer Assisted Intervention Lect Notes Comput Sc* 3216; 43–50.

| Other conference proceeding papers

Bosch, J. G., **M. van Stralen**, M. M. Voormolen, B. J. Krenning, C. T. Lancée, J. H. C. Reiber, A. F. W. van der Steen, and N. de Jong [2005]. Improved spatiotemporal voxel space interpolation for 3D echocardiography with irregular sampling and multibeam fusion. *Proc IEEE Int Ultrason Symp* 2;

— [2006]. Novel spatiotemporal voxel interpolation with multibeam fusion for 3D echocardiography with irregular data distribution. *Proc SPIE Med Imaging* 6147; 61470Q.

Leung, K. Y. E., **M. van Stralen**, M. M. Voormolen, G. van Burken, A. Nemes, F. J. ten Cate, M. L. Geleijnse, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2006a]. Registration of 2D cardiac images to real-time 3D ultrasound volumes for 3D stress echocardiography. *Proc SPIE Med Imaging* 6144; 614418.

Leung, K. Y. E., **M. van Stralen**, G. van Burken, M. M. Voormolen, A. Nemes, F. J. ten Cate, M. L. Geleijnse, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2006b]. Sparse appearance model based registration and segmentation of 3D echocardiographic images. *Proc IEEE Int Ultrason Symp*. 2413–2416.

Leung, K. Y. E., **M. van Stralen**, M. M. Voormolen, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2008]. Improving 3D active appearance model segmentation of the left ventricle with Jacobian tuning. *Proc SPIE Med Imaging* 6914; 69143B.

- Nathanail, K., **M. van Stralen**, C. Prins, F. van den Adel, P. J. French, N. de Jong, A. F. W. van der Steen, and J. G. Bosch [2008, in press]. Rapid 3D transesophageal echocardiography using a fast rotating multiplane transducer. *Proc IEEE Int Ultrason Symp.*
- van Stralen, M.**, J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, C. T. Lancée, N. de Jong, and J. H. C. Reiber [2004]. A semi-automatic endocardial border detection method for the left ventricle in 4D ultrasound data sets. *Proc Computer Assisted Radiology and Surgery Int Congress Series* 1268; 1078–1083.
- van Stralen, M.**, M. M. Voormolen, G. van Burken, B. J. Krenning, R. J. M. van Geuns, E. Angelié, R. J. van der Geest, C. T. Lancée, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2005a]. A novel dynamic programming based semi-automatic endocardial border detection method for 4D cardiac ultrasound. *Proc IEEE Int Ultrason Symp.* 1232–1235.
- van Stralen, M.**, J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, R. J. M. van Geuns, E. Angelié, R. J. van der Geest, C. T. Lancée, N. de Jong, and J. H. C. Reiber [2005c]. Semi-automatic border detection method for left ventricular volume estimation in 4D ultrasound data. *Proc ASCI Conf.* 200–207.
- van Stralen, M.**, J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, R. J. M. van Geuns, E. Angelié, R. J. van der Geest, C. T. Lancée, N. de Jong, and J. H. C. Reiber [2005b]. Semi-automatic border detection method for left ventricular volume estimation in 4D ultrasound data. *Proc SPIE Med Imaging* 5747; 1457–1467.
- van Stralen, M.**, K. Y. E. Leung, M. M. Voormolen, N. de Jong, A. F. W. van der Steen, J. H. C. Reiber, and J. G. Bosch [2007a]. Automatic segmentation of the left ventricle in 3D echocardiography using active appearance models. *Proc IEEE Int Ultrason Symp.* 1480–1483.
- [2007b]. Fully automatic detection of the left ventricular long axis and mitral valve plane in 3D echocardiography. *Proc IEEE Int Ultrason Symp.* 2413–2416.
-

| Submitted papers

- Leung, K. Y. E., **M. van Stralen**, M. G. Danilouchkine, N. de Jong, A. F. W. van der Steen, and J. G. Bosch [2008, submitted]. Motion-guided optical flow tracking for segmenting 4D echocardiograms. *IEEE T Med Imaging*.
- Ma, M., **M. van Stralen**, J. H. C. Reiber, J. G. Bosch, and B. P. F. Lelieveldt [2008, submitted]. Model driven quantification of left ventricular function from sparse single-beat 3D echocardiography. *IEEE T Med Imaging*.

| Abstracts

- Bosch, J. G., **M. van Stralen**, K. Y. E. Leung, M. M. Voormolen, N. de Jong, and A. F. W. van der Steen [2007]. Quantification of left ventricular volume and endocardial wall motion in real-time 3D echocardiography. *Proc First Dutch Conf on Bio-Medical Engineering*. 140.
- Leung, K. Y. E., A. Nemes, G. van Burken, **M. van Stralen**, J. G. Bosch, O. I. Soliman, F. J. ten Cate, and M. L. Geleijnse [2008]. Improving interobserver agreement of three-dimensional stress echocardiography using novel side-by-side viewing software. *Eur J Echocardiog* 9;S1; S160.
- Nemes, A., K. Y. E. Leung, G. van Burken, **M. van Stralen**, J. G. Bosch, O. I. Soliman, B. J. Krenning, W. B. Vletter, F. J. ten Cate, and M. L. Geleijnse [2008]. Side-by-side viewing of anatomically aligned left ventricular segments in three-dimensional stress echocardiography. *Eur Heart J* 29;S1;
- Nevo, S. T., **M. van Stralen**, A. M. Vossepoel, J. H. Reiber, N. de Jong, A. F. W. van der Steen, and J. G. Bosch [2006]. Automated tracking of the mitral valve ring motion in apical echocardiographic images. *Eur J Echocardiog* 7;S1; S120.
- van Stralen, M.**, J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, R. J. M. van Geuns, C. T. Lancée, N. de Jong, and J. H. C. Reiber [2004]. Semi-automatic left ventricular endocardial border detection method for 4D ultrasound data. *Eur J Echocardiog* 5;S1; S57.
- van Stralen, M.**, J. G. Bosch, M. M. Voormolen, G. van Burken, B. J. Krenning, R. J. M. van Geuns, E. Angelié, R. J. van der Geest, C. T. Lancée, N. de Jong, and J.

H. C. Reiber [2005]. Evaluation of automated full cycle left ventricular volume estimation for real-time 3D echo against MRI. Eur J Echocardiog 6;S1; S131.

Dankwoord

Zeker zo aan het einde van de rit, wil ik graag een aantal mensen bedanken, voor hun bijzondere bijdrage aan dit werk. Ik wil alle collega's van het LKEB bedanken voor de prettige start die ik bij jullie heb kunnen maken. Boudewijn, bedankt voor de mooie tijd op verschillende congressen. *Meng, thanks for collaborating on FRU segmentation, and for continuing the PhD student symposia.* Alize bedankt voor het samen opzetten van die symposia.

Ook iedereen bij BME in Rotterdam wil ik bedanken voor de relaxte sfeer, jullie betrokkenheid en de ontspanning binnen en buiten het lab. Gerard, jij hebt een belangrijke rol gespeeld in die prettige werksfeer, in Leiden en Rotterdam, altijd enthousiast voor een nieuw idee en erg meelevend en behulpzaam. Esther van jouw frisse gestructureerde blik op het onderzoek heb ik een hoop geleerd. Net als van de discussies met jou en vooral met Marco. Marco, ik hoop ooit net zo'n trotse vader te worden. Mike, bedankt voor je zeer gewaardeerde uitleg en adviezen. *Folkert, Marcel, Boudewijn, Ossama, Attila and others, thanks for your interest in our research and for being so cooperative. Shelly and Kyriakos, thanks for being such great and inspiring students.*

Verder wil ik mijn collega's bij het ISI bedanken voor het enthousiasme, de motivatie en de tijd die jullie mij hebben gegeven om dit boekje af te ronden. Iedereen van BGR bedankt voor de enthousiaste samenwerking en de kijk in jullie keuken.

Leden van die andere P.C.: Annemieke, Erik, Judith, Loes, Maike, Mieke, Nabil, Ronald, Stephanie en in het bijzonder paranymfen Bart en Stefan, maar ook Hilde, Lars en Marieke, Jeroen en Vera, Sebas, bedankt voor de jullie interesse, steun en vriendschap. Tovonaren, bedankt voor de nodige ontspanning en sportiviteit.

Roel, Wil, Mirjam en Sam, jullie hebben veel voor mij betekend tijdens deze promotie, en zijn altijd erg begaan met alles wat ik doe. Ik waardeer jullie betrokkenheid enorm en de steun die daar vanuit gaat.

Ward, Nienke, Pleun, pap en mam, jullie hebben me altijd gesteund in wat ik wilde doen, of anders wel in het vinden van dat wát ik eigenlijk wilde doen. Zo ben ik op dit punt gekomen, en zonder jullie stimulans, veiligheid, advies, lessen, voorbeelden, fouten en idealen had ik dit nooit gedaan.

Iris, je hebt wat pieken en dalen met me mee beleefd op de weg hier naar toe. Letterlijk, met een hoop geweldige bestemmingen, maar natuurlijk vooral in de figuurlijke zin. Ik ben ontzettend gelukkig met jou en blij dat we niet tot ná mijn promotie gewacht hebben met trouwen. Ontzettend bedankt voor je onvoorwaardelijke liefde en steun!

Marijn van Stralen
December 2008

Curriculum Vitae

Marijn van Stralen was born in Roermond, the Netherlands, on March 15, 1980. He obtained his VWO degree at the Stedelijk Lyceum in Roermond in 1998. In 2003, he earned his Master's Degree in Medical-Technical Computer Science at Utrecht University, by completing his internship at TNO Human Factors in Soesterberg, on analysis of 3D meshes generated by 3D human body scanners.

He started his professional career as a PhD student on automated analysis of 3D echocardiography in 2003. At first in Leiden at the Laboratory for Clinical and Experimental Image Processing (LKEB) at the Leiden University Medical Center and from 2005 at the department of Biomedical Engineering (Thoraxcenter) at the Erasmus Medical Center in Rotterdam, the Netherlands. The results of his research on this subject are summarized in this thesis.

Since 2008, he has taken a position as post-doc researcher on multimodal image registration and segmentation, primarily of pre- and intra-operative brain images at the Image Sciences Institute at the University Medical Center in Utrecht, the Netherlands.

