

Semi-Automatic Segmentation of the Mitral Valve Leaflets on 4D Ultrasound Images

Masterarbeit

zur Erlangung des Grades eines Master of Science (M.Sc.)
im Studiengang Computervisualistik

vorgelegt von

Nils Lichtenberg
nlichtenberg@uni-koblenz.de

Erstgutachter: Prof. Dr.-Ing. Stefan Müller
(Institut für Computervisualistik, AG Computergraphik)
Zweitgutachter: Sandy Engelhardt, MSc.
(Deutsches Krebsforschungszentrum (DKFZ))

Koblenz, im Januar 2015

Erklärung

Ich versichere, dass ich die vorliegende Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Ja Nein

Mit der Einstellung der Arbeit in die Bibliothek bin ich einverstanden.

Der Veröffentlichung dieser Arbeit im Internet stimme ich zu.

.....
(Ort, Datum)

.....
(Unterschrift)

Zusammenfassung

Die Mitralklappe ist eine der vier Herzklappen des Menschen und in der linken Herzkammer zu finden. Ihre Funktion ist es, den Blutfluss vom linken Atrium zum linken Ventrikel zu regeln. Pathologien können zu eingeschränkter Funktionalität der Klappe führen, sodass Blut zurück ins Atrium fließen kann. Patienten, die von einer Fehlfunktion betroffen sind, leiden möglicherweise an Erschöpfung und Schmerzen in der Brust. Die Funktionalität kann chirurgisch wiederhergestellt werden, was meist ein langer und anstrengender Eingriff ist. Eine gründliche Planung ist daher nötig, um eine sichere und effektive Operation zu garantieren. Dies kann durch prä-operative Segmentierungen der Mitralklappe unterstützt werden. Eine post-operative Analyse kann den Erfolg eines Eingriffs feststellen.

Diese Arbeit wird bestehende und neue Ideen zu einem neuen Ansatz kombinieren, der zur (semi-)automatischen Erstellung solcher Mitralmodelle dienen kann. Der manuelle Anteil garantiert ein Modell hoher Qualität, während der automatische Teil dazu beiträgt, wertvolle Arbeitszeit zu sparen. Die Hauptbeiträge des automatischen Algorithmus sind eine ungefähre semantische Trennung der beiden Mitralsegel und ein Optimierungsprozess, der in der Lage ist, eine Koaptations-Linie und -Fläche zwischen den Segeln finden.

Die Methode kann eine vollautomatische Segmentierung der Mitralsegel durchführen, wenn der Annulusring bereits gegeben ist. Die Zwischenschritte dieses Vorgangs werden in eine manuelle Segmentierungsmethode integriert, sodass ein Benutzer den Gesamtprozess beeinflussen kann.

Die Qualität der generierten Mitralmodelle wird durch das Vergleichen mit vollständig manuell erstellten Modellen gemessen. Dies wird zeigen, dass übliche Methoden zur Bestimmung der Qualität einer Segmentierung zu allgemein gefasst sind und nicht ausreichen, um die echte Qualität eines Modells widerspiegeln zu können. Folglich führt diese Arbeit Messungen ein, die in der Lage sind, eine Segmentierung der Mitralklappe detailliert und unter Betracht anatomischer Landmarken bewerten zu können.

Neben der intra-operativen Unterstützung eines Chirurgs, liefert eine segmentierte Mitralklappe weitere Vorteile. Die Möglichkeit, die Anatomie einer Klappe patientenspezifisch aufzunehmen und objektiv zu bewerten, könnte als Grundlage für zukünftige medizinische Forschung in diesem Bereich dienen. Die Automatisierung erlaubt dabei das Bearbeiten großer Datenmengen mit reduzierter Abhängigkeit von Experten. Desweiteren könnten Simulationsmethoden, welche ein segmentiertes Modell als Eingabe nutzen, das Ergebnis einer Operation vorhersagen.

Abstract

The mitral valve is one of the four valves in the human heart. It is located in the left heart chamber and its function is to control the blood flow from the left atrium to the left ventricle. Pathologies can lead to malfunctions of the valve so that blood can flow back to the atrium. Patients with a faulty mitral valve function may suffer from fatigue and chest pain. The functionality can be surgically restored, which is often a long and exhaustive intervention. Thorough planning is necessary to ensure a safe and effective surgery. This can be supported by creating pre-operative segmentations of the mitral valve. A post-operative analysis can determine the success of an intervention.

This work will combine existing and new ideas to propose a new approach to (semi-)automatically create such valve models. The manual part can guarantee a high quality model and reliability, whereas the automatic part contributes to saving valuable labour time. The main contributions of the automatic algorithm are an estimated semantic separation of the two leaflets of the mitral valve and an optimization process that is capable of finding a coaptation-line and -area between the leaflets.

The segmentation method can perform a fully automatic segmentation of the mitral leaflets if the annulus ring is already given. The intermediate steps of this process will be integrated into a manual segmentation method, so that a user can guide the whole procedure.

The quality of the valve models generated by the method proposed in this work will be measured by comparing them to completely manually segmented models. This will show that commonly used methods to measure the quality of a segmentation are too general and do not suffice to reflect the real quality of a model. Consequently the work at hand will introduce a set of measurements that can qualify a mitral valve segmentation in more detail and with respect to anatomical landmarks.

Besides the intra-operative support for a surgeon, a segmented mitral valve provides additional benefits. The ability to patient-specifically obtain and objectively describe the valve anatomy may be the base for future medical research in this field and automation allows to process large data sets with reduced expert dependency. Further, simulation methods, that use the segmented models as input, may predict the outcome of a surgery.

Contents

I Preface	1
1 Introduction	1
1.1 Motivation	1
1.2 Organization of this Work	2
II Background	4
2 Medical Background	4
2.1 Mitral Valve	5
2.2 Mitral Regurgitation	6
2.3 Mitral Valve Repair	7
3 Technical background	7
3.1 Ultrasound Imaging	7
3.2 Segmentation Methods	11
3.2.1 Deformable Models	11
3.2.2 Graph Cut	14
4 Related Work	15
4.1 The Medical image Interaction Toolkit	15
4.2 Mitral Valve Modelling Approaches	16
4.2.1 Learning based Method	16
4.2.2 Active Contour Method	17
4.2.3 Model based Method	20
4.2.4 Other Methods	21
4.2.5 Method Summary	23
4.3 Workstations in Clinical Use	25
III Methods and Evaluation	27
5 Objectives	27
5.1 Benefits from a Segmented Valve	28
5.2 Manual Segmentation	29
5.3 Semi-automatic segmentation	30
6 The Segmentation Pipeline	30
6.1 Manual Segmentation	31
6.2 Pipeline Assembly	36
6.3 Semi-Automatic Pipeline Components	39

6.3.1	Image and Model Alignment	39
6.3.2	Masking the Volume	40
6.3.3	Thin Tissue Detector	40
6.3.4	Maximum Intensity Projection	41
6.3.5	Clustered Graph Cut	42
6.3.6	Leaflet Separation	45
6.3.7	Model Initialization	48
6.3.8	Model Optimization	52
7	Implementation	57
7.1	Mitralyzer Extension	57
7.2	The SeMitralyzer Plugin	60
7.3	Mitral Util	65
8	Experiments and Results	66
8.1	Experiments	66
8.1.1	Magnitudes	67
8.1.2	Process	71
8.2	Results	72
8.2.1	Semi-Automatic Segmentation	73
8.2.2	Manual Segmentation Time Measure	80
8.2.3	Runtime Analysis	81
9	Discussion	82
9.1	Measurements	82
9.2	Manual Modelling and Ground Truth	87
9.3	Semi-Automatic Segmentation	89
IV	Conclusion	90
10	Summary	90
11	Contributions	91
12	Future Work	92
12.1	Model Representation	93
12.2	Shape Prior	93
12.3	Inter- and Intra-user Variability	94
12.4	3D Leaflet Separation	95
12.5	SeMitralyzer Work-flow	96

Part I

Preface

1 Introduction

1.1 Motivation

According to the American Heart Association 2% to 3% of the US american population suffers from mitral valve prolapse (MVP) [1] which is a heart disease where the mitral valve (MV) does not close completely during heart contraction. This leads to a portion of the blood flowing back from the ventricle into the atrium, which is called mitral regurgitation (MR). Diseases like this are generally called mitral insufficiency (MI). Typical symptoms are fatigue, chest pain and dizziness. The malfunction can be resolved by replacing or reconstructing the valve.

An article published in 2007 by the U.S. News and World Report [2] states that only 50% of the treated patients got a mitral valve reconstruction (MVR). The other half got their mitral valves replaced completely. The disadvantage of a full replacement is that either, a tissue valve has to be replaced again about 10-12 years later or, a mechanical valve requires a life long use of blood thinners. Another article published in 2013 by the UHZ in Freiburg [3] depicts present numbers, estimating 11,2% of the common population suffering from a medium MI and 70-90% of the severe cases being treated with a MVR. This trend leads to the assumption that surgeons will more and more turn away from mitral valve replacements and focus on MVR instead.

A way to support MVR with respect to the increasing number of surgeries in this field, is making surgery planning easier and more sophisticated. The latest achievements are improved 3D echocardiography techniques [4]. Especially the transesophageal echocardiography (TEE) allows for high spatial resolution and anatomical detail. Preoperative knowledge about the MV anatomy can have a positive effect on optimizing the surgeon's strategy which may then lead to a better result of the surgery and a more sustainably reconstructed valve.

Present research deals with the recorded TEE images and tries to segment the MV and extract parameters that can describe the extent of a possible MI. An example is displayed in figure 1. This image shows the result of a semi-automatic segmentation of the MV done by Pouch et al. [5]. The model has been quantified by measuring the leaflet thickness, which may unveil anomalies in the MV anatomy.

The disadvantage of TEE imaging is that it can only be applied to anaesthetized patients, meaning that the images are usually acquired right before the surgery. To not disturb the common execution of the surgery, it is necessary that an application requires minimal interaction by a surgeon to obtain the desired data.

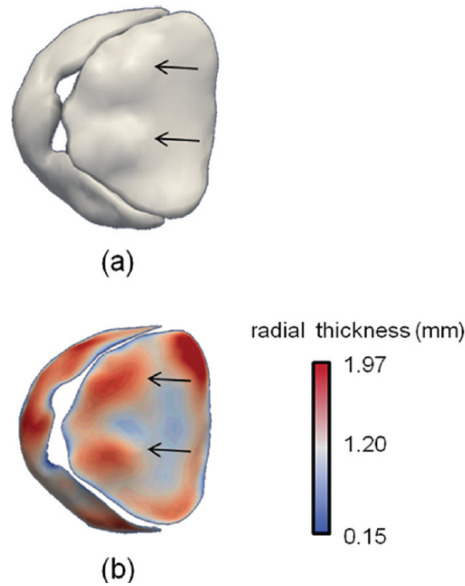


Figure 1: Example of MV segmentation and measurement of leaflet thickness [5].

Consequently this work will develop and implement a semi-automatic application to segment and quantify the MV. The model representation used for the segmentation must be structured in a way so that clinically relevant measurements can be done after the segmentation. Very important to measure the quality of a semi-automatic segmentation is a ground truth segmentation done by an expert. Because of that, a fully manual segmentation method, that delivers reliable results, has to be implemented as well. Automatic steps should then be integrated into the manual method to obtain a semi-automatic work-flow. This requires that the model representation can be used for manual as well as for automatic segmentation. Finally, the results obtained in this work will be discussed, which includes the comparison to other present methods.

1.2 Organization of this Work

Section 2 will describe the medical background of the MV and MVR in detail. The following sections in 3 will deal with TEE imaging and a brief introduction to segmentation methods. The related work section 4 introduces the Medical Image Interaction Toolkit (MITK) (4.1) and the work of other research groups in the field of MV segmentation (4.2). Two clinical

workstations that can be used for MV quantification are presented as well (4.3).

Based on the information discussed to that point, the final objectives for this work are derived in section 5, followed by section 6 and 7 which are about the theory and implementation to achieve those objectives. The experiments and results can be found in section 8, where new magnitudes to measure the segmentation quality of a MV are presented. Section 9 will deal with the discussion of the results obtained during the experiments. The accomplished goals and objectives of this work will then be summarized in section 11 of the final chapter. Finally, section 12 talks about ideas that could be referred to in further development of the methods proposed in this work.

Part II

Background

2 Medical Background

The content of this chapter is based on [6, 7].

The MV is one out of four valves that can be found in the human heart. Its function is to control the blood flow from the left atrium to the left ventricle. The heart anatomy is illustrated in figure 2. In detail, oxygenated blood runs from the lungs into the left atrium. Then, when the heart relaxes, the valve opens and blood flows into the left ventricle. This state is called *diastole*. When the heart contracts, the valve closes and blood runs into the body circulation system. This state is called *systole*. If the valve does not close properly and blood leaks from the left ventricle back into the left atrium, the patient suffers from MR.




Figure 2: Anatomy of the heart (image taken from Wikipedia).

2.1 Mitral Valve

The MV is made up from four major components. The *mitral annulus*, the *leaflets* and the *chordae tendineae* that are attached to the *papillary muscles*. Figure 3 shows the MV from the left atrium point of view. The *anterior-* and *posterior leaflet* are embedded in the *mitral annulus*. The annulus is a ring-like structure that is more rigid than the surrounding tissue. It is crucial for stabilizing the leaflets.

The leaflets touch at the coaptation line during systole, closing the valve and forming a coaptation zone (see figure 3). The end points of the coaptation line are called commissure points.

The papillary muscles are located in the left ventricle. They are connected to the leaflets by strings called chordae tendineae. Their task is to restrain the moveability of the leaflets, so they do not open into the left atrium during heart contraction. The leaflets can also be divided into three segments each. The segments arise from the fact that the posterior leaflet has three clefts that visually divide the leaflet into three independent parts. The three segments on the anterior leaflet are separated according to the posterior segments. This classification scheme was proposed by Carpentier [8].

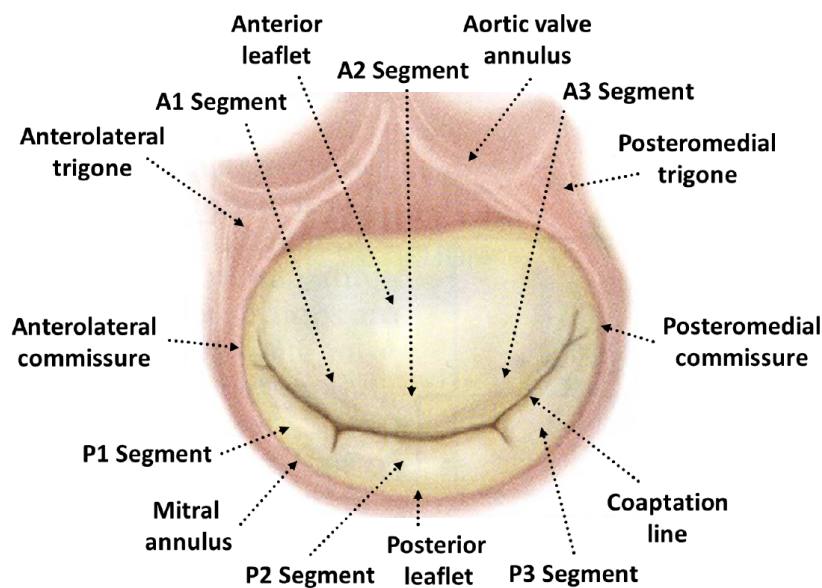


Figure 3: Anatomy of the MV (image taken from [8]).

The following picture shows an exercised leaflet. Even in this optimal view it is hard to find an exact distinction of the leaflet and the chordae.



Figure 4: Image taken from themitralvalve.org, with permission to print by Farzan Filsoufi.

2.2 Mitral Regurgitation

MR is a heart disease that occurs when the two leaflets do not close completely during systole. In this case blood can flow from the left ventricle back to the left atrium. MR can appear if the annulus dilates due to an enlargement of the left ventricle (*mitral annulus dilatation*) or if one or more chordae are ruptured or are too long, which leads to MV prolapse. Rheumatic fever can lead to a thickening of the leaflets, causing decreased leaflet mobility.

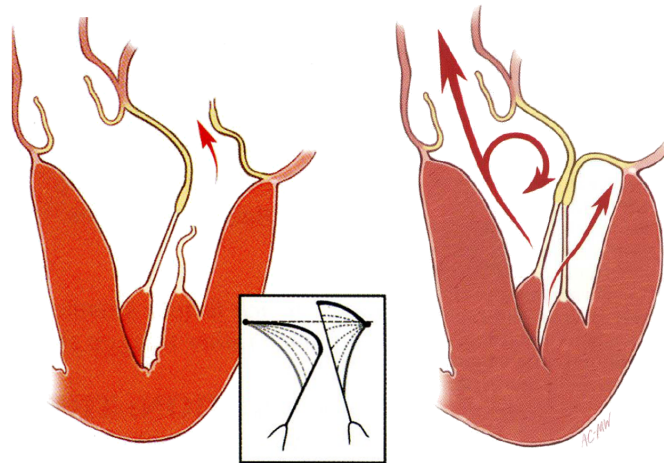


Figure 5: Left: MV prolapse. The chordae have ruptured and cannot restrain the leaflet movement to close the valve properly and blood runs back into the atrium. Right: Healthy valve (image taken from [8]).

Figure 5 shows a ruptured chordae, with the effect of a sudden increase of leakage. If the leakage develops slowly over time the heart tries to compen-

sate the leakage by increasing the ventricle volume. This eventually leads to even more leakage and causes breathlessness or fatigue of the affected person. Severe cases can cause a heart attack, making a surgery indispensable.

2.3 Mitral Valve Repair

A way to treat MR is to attempt reconstructing the valve. In this intervention the original valve tissue is used to restore the valve's functionality. This can be achieved by resectioning the leaflets and/or deforming the annulus with an annuloplasty ring. As described in the work of [6] an annuloplasty ring is used to support the annulus' ability to stabilize the valve structure. It can also be used to resize the valve diameter and thus help the leaflets to close completely again. If a prolapse exists the leaflet can be resectioned to fit. Unfortunately this is usually only possible for the posterior leaflet, according to [3]. Alternatively, artificial strings called neo-chordae can be implanted to restrain the leaflet motion. This is applicable for both leaflets and preserves all leaflet segments which leads to a good coaptation of the leaflets [3]. Figure 6 shows a picture taken by an endoscope during a MV repair surgery.




Figure 6: Atrial view on the MV. The sutures are used to fix an annuloplasty ring to the annulus. The approximate leaflets are highlighted green (posterior) and red (anterior).

3 Technical background

3.1 Ultrasound Imaging

The basic concept of ultrasound (US) imaging is the processing of reflected ultrasound waves into 2D images. A transducer emits the waves into the patient's body and receives the reflection. Different materials or different kinds of tissue vary in their reflectance properties. For example, blood reduces the sound pressure of the reflected waves by up to 120 dB, muscle tissue by up to 50 dB and bone by 20 dB [9]. The received signals can then be

processed to 2D images. Practical aspects about ultrasound are that, on the one hand, it is a cheap, portable, and non-ionizing imaging modality, capable of capturing the fast moving valve structures [10]. On the other hand its points of view are restricted by the patient's body and signal depth. Also, results are highly dependent on the operator's skills.

The first approaches to achieve 3D images based on ultrasound were reconstructions based on a set of 2D images [11]. The disadvantage is, that the content of these 2D images has to be free of motion over time, or motion has to be compensated, to achieve good 3D results. To circumvent this problem, real-time 3D ultrasound probes are equipped with a grid of transducer elements. This enables imaging of a whole volume in a single time step.

Two different techniques are available to obtain real-time 3D US images. The transthoracic echocardiography (TTE), where the images are taken from outside the body and the transesophageal echocardiography (TEE), where the transducer probe is placed inside the esophagus of the patient. As stated before, bone reflects a high portion of the incoming waves. During TTE image acquisition, bones may obstruct an optimal view on the MV [6]. TEE can achieve better images, since the US probe is located inside the body. That makes it possible to move very close to the patient's heart and avoid the view to be occluded by the vertebrae. The results are images of higher quality and higher spatial resolution as stated in [4]. The higher spatial resolution is given by the fact that the TEE probe can be positioned closer to the heart. Figure 7 shows how the image is acquired in a cone-like shape. The slices that the volume is composed of intersect at the probe, so that regions closer to the probe are imaged with a higher spatial resolution. Since good image quality is crucial for this work, TEE images will be used here.

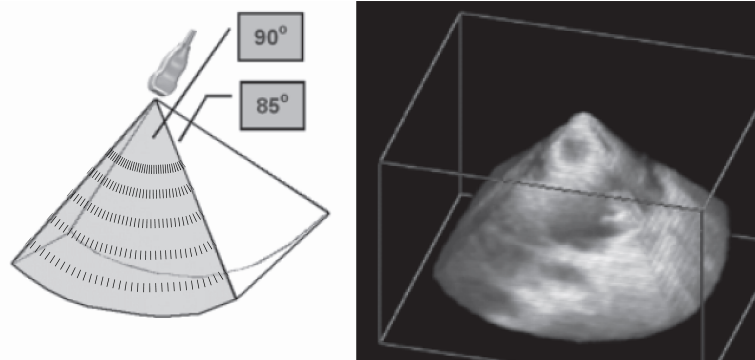


Figure 7: Full volume obtained by a real-time 3D probe (image taken from [11] and edited).

Figure 8 shows how close the TEE probe can be placed to the heart. The esophagus is not shown in this image to give view on the probe. Different positions of the probe can achieve good views for different purposes, the example is in the four-chamber view which is used to show as much of the heart anatomy as possible. The 3D volume has been sliced horizontally for this purpose, but it is also possible to cut the volume in arbitrary slices.

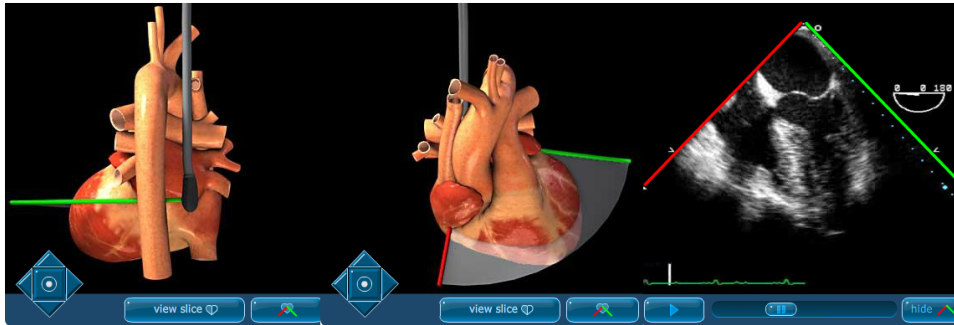


Figure 8: Visualization of a TEE probe placed next to the heart and 2D image slice resulting from 3D acquisition (image taken from the Perioperative Interactive Education (PIE) group at the Toronto General Hospital (<http://pie.med.utoronto.ca>)).




Figure 9: Low signal (A), Shadows (B), Static artifacts (C) and bad 3D stitching (D) (image taken from [6]).

A general problem with ultrasound images is their sensibility to image artifacts. Four common cases are described in [6] and will be shortly men-

tioned here. More kinds of artifacts can be found in [9]. Figure 9A shows a low or missing return signal, shadows as in figure 9B occur if the signal hits a strongly reflecting material and no signal reaches through to deeper regions. Errors in the transducer can cause static artifacts shown in figure 9C. When reconstructing images from a set of images the stitching might fail and result in an image with shifted regions as in figure 9D.

Drop-outs in the US signal can sometimes result in apparently plausible images, that do not fit to the actual anatomy. Figure 10 shows the valve during systole. It looks like there is a gap between the anterior and posterior leaflet as during diastole. But a manual segmentation, considering neighbouring slices and time steps, revealed that the leaflets actually coapt at another location. This error was caused by artifacts during image acquisition.




Figure 10: Difficulty of distinguishing leaflets. Red arrow: Apparent leaflet ends. Green arrow: Actual leaflet ends and coaptation.

The images in figure 11 show further imaging difficulties that arise during a cardiac cycle. There is no intensity based boundary from the leaflets to the remaining tissue. The posterior leaflet is often pressed against the ventricle wall during diastole and there is no visual demarcation between the leaflets when they coapt.




Figure 11: Yellow arrows: Annulus. Green arrows: Posterior leaflet tip. Pink arrow: Coaptation. Image taken from [12].

The images used in this work are obtained by a Philips workstation. The workstation represents the images with polar coordinates but when exporting the images to make them accessible via MITK, the coordinates are

stored in Cartesian system coordinates. Figure 12 illustrated the export effect. The exported image contains far less detail and edges are blurry. This makes image processing outside the workstation more difficult.




Figure 12: Left: Original image. Right: Exported image (image taken from [6]).

3.2 Segmentation Methods

This section will introduce segmentation methods that are common in clinical context. Their general advantage is, that they do not operate merely pixel-based. Instead, they can use global information to overcome artifacts and irregularities in an image.

3.2.1 Deformable Models

Deformable Models (DM) have been researched in medical context for a long time. Artifacts in medical image data made it necessary to go from low-level, simply pixel or voxel based, image processing to model based processing to obtain more reliable segmentation results. Model based image processing has the ability to overcome obstacles such as artifacts and discontinuities in image data and furthermore to integrate a priori knowledge about size, shape and location of the searched structures [13]. Prior knowledge in form of a shape prior model (SP) can be modelled manually from a set of images as training data. According to [14] the use of a prior model has several advantages:

- The same algorithm can be applied to many different problems, merely by presenting different training examples.
- Expert knowledge can be captured in the system in the annotation of the training examples.
- The models give a compact representation of allowable variation, but are specific enough not to allow arbitrary variation different from that seen in the training set.

The SP can be an average model taken from the training set as a first guess to start segmentation from or it can supply additional information such as variability of the training set. The segmentation algorithm can try to mimic this variability as guidance and to restrain the deformation of the model [14].

To create an SP an expert has to segment a set of training images by specifying landmarks that describe the model. Suitable landmarks are found at corners or junctions of object boundaries and biologically significant locations that can be easily detected. The defined points can then be used to define a model, using an appropriate representation. For instance, consecutively defined points could be connected to a contour via a spline representation.




Figure 13: Example of suitable model landmarks (image taken from [14]).

Figure 13 shows an example segmentation and annotation. The very significant landmarks (black dots) are too few to delineate the actual structure, so that additional landmarks are added along the contour.

The drawback of SPs is, that a large number of data sets is required to build a good SP. An SP that is used for (semi-)automatic 4D segmentation has to cover both, possible initial poses of a model with respect to specified landmarks and different configurations over time to model the time-based movement of the model.

A way to optimize a deformable model are *active contours*. This is a segmentation method where the optimal position of a contour is determined by an energy term that consists of an *internal force* and an *external force*. The energy term has to be minimized for an optimal solution. The internal force is the input contour's urge to keep its initial shape and the external force is the attempt of the image to adjust the contour landmarks towards detected image landmarks. The technique has been first proposed by Kass et al. in [15], termed *snakes*. They define the energy function of a snake as follows:

$$E_{snake}^* = \int_0^1 E_{int}(v(s)) + E_{ext}(v(s)) + E_{con}(v(s)) ds$$

where $v(s)$ is a parameter representation of the contour, E_{int} is the internal force, E_{ext} is the external force and E_{con} can be used to apply external constraints.




Figure 14: Active Contour example (image taken from [16]).

Figure 14 shows a contour in a white image region next to a gray image region. The gradient magnitude is computed along the normal of an example parametric point. The external energy term might penalty regions of low gradient magnitude. In this case the gradient information could be used to move the point towards a location where high gradients can be found. At the same time, the internal force might try to keep that point in its location.

This approach can also be applied to 3D use cases. For instance, points that are connected to a 3D surface, instead of a 2D contour, can be optimized with the same methods within a 3D image. Active contours are sometimes referred to as active shapes or active surfaces if used in a 3D context. Also, the energy term can be of any kind. It is not bound to be composited of an external image force and an internal force that tries to keep the surface or contour in its initial shape. The force has to be constructed individually to fit its purpose.

3.2.2 Graph Cut

Part of this work will use a min-cut/max-flow graph cut to label image regions, so the basic idea behind this method will be briefly described in this section and is based on [17].

When using graph cuts in computer vision, the graphs usually represent the same grid as the pixel or voxel grid in the image. Nodes are then connected so that they have the same neighbourhood as the pixels/voxels. The edge weights or capacity between the nodes can be based on pixel intensities or any term based on the information in the image. Some nodes have to be connected to terminal nodes. There have to be as many terminal nodes as desired image labels. When using two labels, the terminals are usually called *sink* and *source*. In this case, the graph can be seen as a pipe-system and the edge weights refer to the maximum capacity of water that can flow from one node to another. The graph cut will then find the edges that are saturated, when the maximum amount of water flows from source to sink. The saturated edges form a boundary between the two terminals that separates the graph into two disjoint subsets. The subset membership of each node can be transferred to the original image pixels, creating two labelled image regions.

An example is given in figure 15. Assume that each node represents a pixel and the node name is its intensity. Each pixel is connected to its neighbours and the edge capacity is the sum of the neighbouring nodes' intensity. The capacity to the terminal nodes is infinite in this example. The edges between the high-intensity and low-intensity pixels are first saturated. The graph can now be separated into two sets by cutting through the saturated edges.




Figure 15: Example graph cut with source s and sink t . The dashed line cuts the two saturated edges.

The fundamental statement about min-cut/max-flow problems is, that the capacity of the saturated edges determines the maximum flow (*max-flow*) in a graph. At the same time, the set of saturated edges can be used to cut

the graph into two regions at minimum cost (*min-cut*). Thus, the max-flow and min-cut represent an equal solution to the same problem.

4 Related Work

4.1 The Medical image Interaction Toolkit

The Medical Imaging Interaction Toolkit (MITK) is a framework for medically based applications [18]. The software is open source and developed by the Department of Medical and Biological Informatics of the German Cancer Research Center (DKFZ). Plug-ins for special purposes can be implemented separately for the MITK-Workbench, which is the base stand alone application supported by MITK. The works of Seitel [19] and Graser [6] were also implemented on MITK base and developed the Mitralyzer plug-in (see figure 16). The work at hand is based on their work.

The Mitralyzer is a plug-in for MITK that allows the manual and semi-automatic segmentation of the mitral annulus on 4D TEE images over a whole cardiac cycle. Quantification of the resulting annulus shape is also supplied. The user only has to set the commissure points for some key time steps to initialize the segmentation. Further adjustment of the annulus shape can then be made either manually or by an optimization process based on active contours.

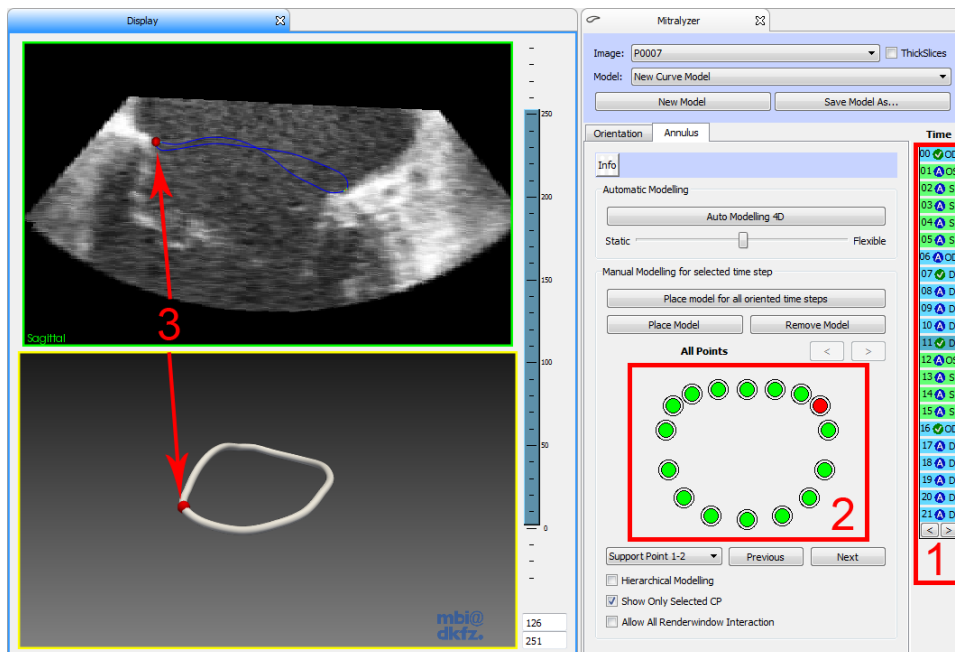


Figure 16: Original Mitralyzer presented in [6]. Time step selection (1), annulus scheme (2) and selected annulus point in 2D and 3D (3).

Figure 16 shows the segmentation view of the Mitralyzer. The modelling can be done for every time step and the curve that represents the annulus is made up by 16 moveable support points. The points can be selected in the annulus scheme on the right hand side. For each point, the user is presented a corresponding slice through the image volume. The slice contains the selected point the, annulus center and is perpendicular to the annulus plane. The user can then adjust the point in the left top window to the anatomically correct annulus position. Alternatively the segmentation can be done automatically. An active contour approach then tries to minimize a cost function that considers a mean shape prior created from a set of manually segmented valves.

4.2 Mitral Valve Modelling Approaches

Modelling the MV from 4D-TEE ultrasound images is a relatively new field of research. First work has been done 2008 by Ionasec et al. in [20], proposing a model driven segmentation approach for the aortic valve using several anatomical landmarks. In a subsequent paper [21] from 2010 they refined their method and also applied it on the MV. Other popular work was done by Schneider [22], relying more on basic image processing and by Pouch [5], utilizing a complex model representation. The following sections will refer to an error or accuracy of the presented methods. If not commented otherwise, this error is the root mean squares error (RMS) between a (semi-)automatically segmented model and a manually segmented model by an expert.

4.2.1 Learning based Method

The process of Ionasec's work is completely automatic and divided into three steps (see figure 17 and [21]):

- **Global Location and Rigid Motion Model:** The rough location and orientation (e.g. location, scale and rotation of the bounding box) of both valves is determined by a learning based approach called Marginal Space Learning (MSL). MSL splits high dimensional spaces into spaces of lower dimensionality and searches these spaces, resulting in lower computation times. Eventually the low-dimension results are combined to the original space. The localization also takes the spatial movement over time into account, as well as Haar-[23] and steerable features [24] and probabilistic boosting trees [25] for clustering hypotheses using the detected features.
- **Nonrigid Landmark Motion Model:** Both valves are represented by landmarks. The landmarks are used to represent the time dependent movement of the valves during a cardiac cycle. Possible movement

trajectories are learned for x , y and z direction separately and form a trajectory spectrum restricting the behaviour of the valve apparatus.

- **Non-Rigid Shape Estimation:** Based on the location of the landmarks and motion prior knowledge a boundary detector delineates the surface of the valves and transfers the detected shape to a mesh. At first cardiac key phases are estimated using a time step at end systole and end diastole, then the remaining time steps are detected using a motion prior.

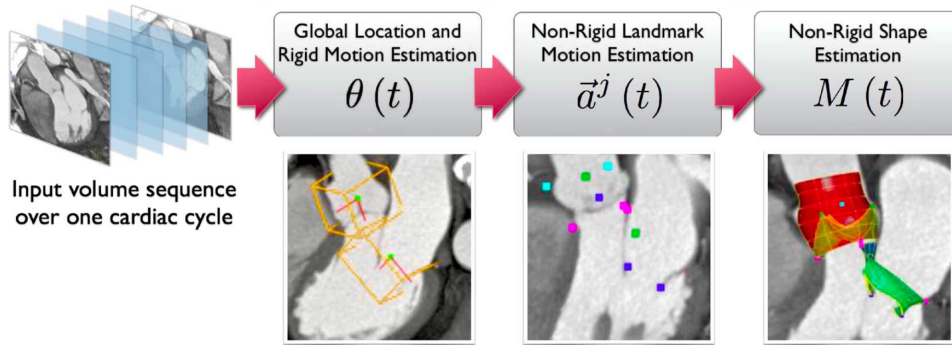


Figure 17: Localization and refinement of aortic (red) and mitral (green) valve apparatus in three steps (image taken from [21]).

The results of Ionasec’s method have been compared to manual measurements done by experts and yield a precision of 1.54 mm on TEE images. This is a very good outcome since this is in the range of ultrasound image precision. Yet, the results have to be read with caution. Even though they used a collection of 65 TEE and 59 CT datasets the learning based elements of the method may be overly fitted to the available training data volumes and thus deliver good results for similar volumes but may fail on data representing other pathologies. Also, the method may be missing leaflet detail due to mainly fitting the model to a few landmarks.

4.2.2 Active Contour Method

Schneider et al. have published several papers on the semi-automatic segmentation of the mitral annulus and leaflets [10, 26, 22, 27]. Their work is fundamentally based on a *Thin Tissue Detector* (TTD) that was designed to detect regions of ultrasound images that may represent thin tissue. A thin tissue detection like this was first found in [28]. For each voxel in an image the TTD computes a value that describes the voxel’s tendency to be located within such an area. These values are used to establish a graph (see [17] and section 3.2.2) that represents the image, with each node’s weight based on the TTD value of the corresponding voxel. This graph is cut in a min-cut/max-flow manner. The location of the cut edges in the image can then

be used to form a surface which delineates the leaflets. Further refinement is done to acquire the annulus location and precise leaflet surface. Figure 18 shows an example of a TTD result image. The TTD algorithm will be explained in detail in section 6.3.3. In [10, 26] the TTD value is used to find the mitral annulus. Only a single user defined point that should be located close to the center of the valve and a predefined search space around that point is required to initialize the segmentation of the annulus.




Figure 18: Example result of the Thin Tissue Detector, MV annulus (MVA), left atrium (LA), left ventricle (LV), and ventricular septum (VS). The dark regions in the right image are regions of thin tissue detected by the TTD (image taken from [10]).

Since the work at hand focusses on leaflet segmentation, only the affiliated parts found in [22, 27] will be explained here.

The leaflet segmentation also utilizes the TTD, using its grey-values to define node weights in a graph. The graph is constructed based on a search space defined by an axis and points on the annulus. Figure 19 shows a schematic of the construction. The image on the left shows planes (dashed lines) that are rotated about the axis in angular steps $\Delta\theta_p$. For each plane (right image) arcs are constructed with radius R increasing with ΔR . The arcs are centered at the cut point of the annulus and the plane. On each arc, nodes are distributed in an angular distance $\Delta\theta_s$. The graph is then created by connecting a node at (θ_p, θ_s, R) with neighbours at $(\theta_p \pm \Delta\theta_p, \theta_s \pm \Delta\theta_s, R \pm \Delta R)$. The edges among neighbours are weighted based on the TTD values and designed so that a min-cut/max-flow graph cut (see [17] and section 3.2.2) is likely to cut the graph in a plane along the leaflet location. Since the graph cut produces a closed surface, leaflets of open valves are connected through the blood pool. This section has to be trimmed out of the plane. After trimming the plane an active surface algorithm refines the segmentation and controls curvature of the leaflets. Compared to man-

ual delineation of the leaflets by an expert, the algorithm achieved a mean accuracy of $0.76 \pm 0.65\text{mm}$ on four data sets and 30 time steps.

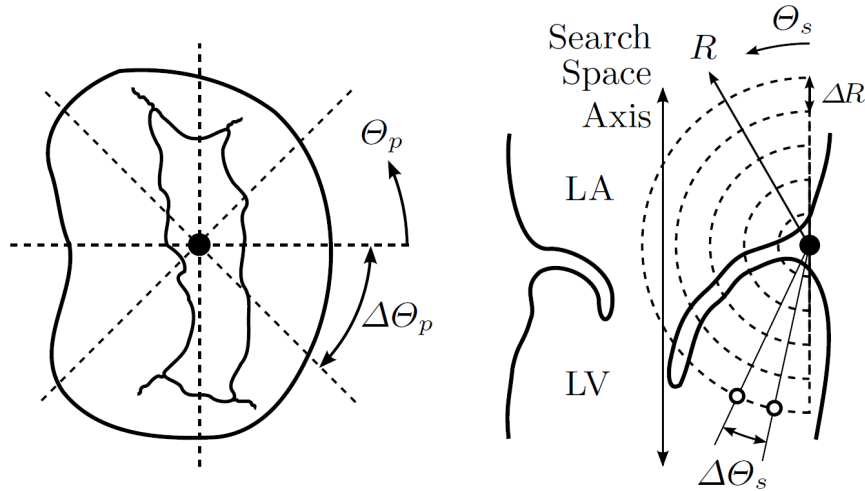


Figure 19: Left: Atrial view of an open MV showing search planes oriented about the search axis. Right: Cut-plane view of an open MV showing the search arc and sample point configurations. (LA - left atrium LV - left ventricle, image taken from [22]).

The most recent publication by Schneider [27] is based on a single time step segmentation. The resulting mesh is then evolved over time as an active surface to fit the remaining time steps. The force equation to evolve the mesh consists of several parameters, including physical constraints of the leaflets. The parameters are:

- **Artificial strain** to maintain original leaflet length. Leaflets can stretch up to 35% in radial direction and 20% in circumferential direction.
- **Bending force** to straighten the leaflets. The ultra sound data may insinuate a folding due to imaging artifacts.
- **Collision force** to avoid leaflet intersection and support the definition of a coaptation line
- **Tether force** to simulate the constricting effect of the chordae on the leaflets

The method is able to segment a sequence from a open to closed valve with a mean accuracy of 1.06 mm. Unfortunately, it has only been evaluated on two data sets which is not enough to make definite statement on the performance.

4.2.3 Model based Method

A third segmentation approach is presented by Pouch et al [5]. Their method requires less image preprocessing but more user interaction. The user has to define a 3D region of interest and then has to delineate the annulus and coaptation line in a 2D projection of the ROI. Figure 20 shows the three initialization steps. At first the user sets two points that define axial planes resulting in a ROI containing the MV. A maximum intensity projection is then made in axial direction (see middle and left image). On the projection image, the user now defines several points to describe the contour of the annulus and the anterior leaflet. Two masks can be created from the contours, including the posterior and anterior leaflet respectively. The original location of the maximum intensity voxels within the mask now form a point cloud that can be dilated to obtain a ROI which includes the leaflets. This ROI is then treated as an active contour to find a first estimation of the MV.

For further processing a template shape model is required. This model is generated based on a sample segmentation that has run through the steps described above. From the dilated maximum intensity voxels a 3D skeletonization is formed. The vertices of the skeleton are used to create a Voronoi diagram. Voxels that do not withstand specific constraints get pruned in order to enforce a smooth shape of the diagram.




Figure 20: Left: Two points define the ROI containing the valve. Middle: The user sets points to outline the annulus. Right: the user outlines the anterior leaflet, defining the coaptation (image taken from [5]).

The plane given by the skeleton of the 3D Voronoi diagram is the base to form a 3D volume representation of the leaflets by the concept of *continuous medial representation* (cm-rep). This concept defines a 3D shape by a skeleton plane that gets "inflated" to a certain thickness. The generated cm-rep shape prior is used to form a mesh with proper vertex topology. This mesh is then fitted to the initial segmentation of the active contour. Figure 21 shows the two models and the final segmentation. Comparing the accuracy of the semi-automatic segmentation to a manual segmentation has resulted in a

mean error of 1.3 ± 0.7 mm on 14 data sets, using one midsystolic time step of each data set.




Figure 21: Left: Initial segmentation. Middle: Shape prior. Right: Final segmentation (image taken from [5]).

4.2.4 Other Methods

This section will briefly introduce three additional methods. They are not described at the same detail as the others, because they follow different basic concepts. The first method works on 2D images only, and the second method uses a manually segmented model as input, so their results are not comparable to the methods described above, whose task it is to find the valve anatomy in general. The third method is a physically-driven extension to the work of Ionasec [21].

An approach that was designed as a preprocessing for real-time segmentation was published by Martin et al. [29] in 2006 for 2D US images. The segmentation has to be initialized with a manual segmentation for one time step, delineating the leaflet with one curve and the cardiac muscle with another curve. The intersection of both curves is approximately a point on the annulus (see figure 22).




Figure 22: Snakes for the cardiac muscle and leaflet (image taken from [29]).

It is assumed that the cardiac muscle performs an almost rigid, translational movement over time with minor rotational components. Based on

that, the contour is tracked from the initial segmentation over consecutive time steps. That is done by estimating the new location and then refining it using active contours. The same is done for the leaflet contour, based on the now given intersection point to the cardiac contour. An energy term restricts stretching of the contour to maintain the length of the initial segmentation. Figure 23 shows an example segmentation.

The tracking is done for several cardiac cycles and in near-real time. The authors planned to use the information extracted from this segmentation to accelerate the segmentation on the same patient at a later time. Since the anatomy and movement of the valve should still be the same, it was assumed to be possible to track the leaflets in real time during surgery.

Unfortunately, a publication about this next step could not be found. A problem faced could have been, since the algorithm works on 2D+T images, that the cardiologist was not able to find the same position for the ultrasound probe to obtain the same view on the MV. In this case the pre-processed segmentation would be futile.




Figure 23: Example segmentation (image taken from [29]).

Another approach that focussed less on automation and more on precision can be found in [30]. This work is also based on a TTD. Instead of Schneider's method, the one proposed in [28] is used. An automatic segmentation of a valve during diastole is presented to an expert. The expert refines this segmentation. The manually created model is then prepared for simulation and during the simulation process the closed valve is estimated, based on the TTD values and a physical representation of the valve. Finally, a manual segmentation of the valve in closed state is used to measure the prediction error of the simulation. Figure 24 shows a manually segmented open valve and an estimated closed valve.




Figure 24: Left: Expert refined open valve state. Right: Predicted closed valve state (image taken from [30]).

The most interesting fact about Burlina’s method is, that it can be used to predict the outcome of a surgery. Once the manually segmented model is available, it can be modified to, for instance, simulate the resection of a leaflet. Of course there is no way to measure the quality of this prediction, only the prediction error computed for the original closed valve can be used for an estimated confidence.

A work that is partly based on Ionasec’s [21] was presented by Voigt et al. [31]. They augmented the data-driven approach by Ionasec [21] with a physically-driven refinement. The data-driven method predicts vertex locations with a certain confidence. This confidence is then considered while further refining the model physically-based with information from other time steps. Assumptions about the leaflet tissue properties are made to be able to model them physically. Their method achieves a mean point-to-mesh error of 3.11 mm for the leaflets and mean Hausdorff distance (see section 8.1) from a model to its ground truth of 5.57 mm. Compared to the data-driven approach, the mean point-to-mesh error is larger, whereas the Hausdorff distance is smaller. Voigt interprets this as a hint that the physically-driven method is more robust.

4.2.5 Method Summary

This section will briefly summarize the three mainly presented segmentation methods for a quick overview. The results of these methods will later be compared to the results obtained in the work at hand. The methods by Martin and Burlina are neglected for the reason stated in section 4.2.4.

A direct comparison can only be made in terms of (semi-) automatic to manual segmentation differences, since that is the only measurement that is supplied for all three approaches. Among that, the complexity of the

applied algorithms, the required user interaction and the clinical relevance are listed.

Ionasec's [21] method uses a wide range of sophisticated utilities to achieve a fully automatic segmentation. For a completely comprehensive reading, one has to have background knowledge in the fields of machine learning, tracking, shape models and techniques that partly have only recently been developed such as marginal space learning.

Schneider [22, 10] uses their TTD, which is composed of basic image processing operations and a graph cut, as well as classic active surface optimization, utilizing several components to form a force term. The required user interaction is very simple and could also be done by a non professional.

The method proposed by **Pouch** [5] is fundamentally based on a specific model representation. While the initial segmentation is made using active contours and therefore similar to Schneider's final step, the model representation allows for modelling of a volumetric valve. Leaflet thickness can be modelled that way, which separates this approach from the others.

The clinical relevance of all three methods is quite similar, but based on different aspects. Ionasec states that processing one volume takes 4.8 seconds what would make it a good method for intra-surgery use. Schneider achieves a more detailed model for the leaflets and coaptation line but gives no information about processing time and may therefore be useful for pre- and post-operative measurements. The same counts for Pouch whose method requires expert user initialization but delivers a volumetric model. An overview, summarizing the key data of each method can be found in table 1.

The presented methods were all validated by comparing the automatic or semi-automatic segmentations with manual segmentations. But how accurate are manual segmentations? Siefert et al. [32] have investigated the accuracy of manual segmentations by using a heart simulator with real porcine MVs. The valves were imaged by an ultrasound probe in an active, moving state within the simulator and segmented based on the resulting image data. Visual markers on the valve were also tracked to create an exact ground truth model of the valve. The mean error between the ground truth and the manual segmentation was 0.34 ± 0.30 mm, which is greater than image precision but lower than an error necessary to lead the surgeon to make a wrong decision.

Author	Available Data	Mean Error (mm)	Required User Interaction	Complexity
Ionasec	1516 time volumes	1.54 ± 1.17	none	high
Schneider	4 hand selected time volumes	0.76 ± 0.65	simple	medium
Pouch	14 hand selected time volumes	1.3 ± 0.7	expert	medium-high

Table 1: Quick overview over the works of Schneider [22, 10], Pouch [5] and Ionasec [21].

Additionally, the Hausdorff distance achieved by Voigt has to be kept in mind, since it is the only of the proposed papers that mentions this magnitude, which is reported with an average of 5.57 mm.

4.3 Workstations in Clinical Use

There are two broadly used clinical workstations that offer segmentations of the MV: QLab by Philipps [33] and 4D MV-Assessment by TomTec [34]. Both programs are designed to segment and quantify the MV to support clinicians during treatment and decision making. The TomTec software supports segmentation on 4D images over systolic time steps (closing of the valve) whereas QLab only supports segmentation for one end-systolic time step.

A review on an earlier version of TomTec states that significant manual correction of the semi-automatically segmented leaflets was required [35]. The key features are the visualization and quantification of prolapse topology, regurgitant orifice, coaptation line and jet origin.




Figure 25: Example visualization of TomTec MV (image taken from [34]).

The software was not available for testing during the work on this thesis and no further statements on quality and user interaction can be made.

QLab focusses on the segmentation and quantification of one end-systolic time step. This is the moment right before the valve starts to open which is

important for clinicians and for measuring MV pathologies. The segmentation is a mixed procedure of user defined landmarks and semi-automatic segmentation. The user is guided step-by-step through the segmentation by the software. Figure 26 shows the leaflet segmentation step. The program proposes an automatic delineation and the user can refine the support points if desired.




Figure 26: Example segmentation in QLab. The points were proposed automatically and can be refined manually.

This combination of user initialized segmentation and manual refinement makes it possible to detailedly segment one time step in less than 5 minutes. This time was estimated after a short user test. Anyway, problems occurred during the process, since the slices presented to the user are not always optimal and lead to misinformation. In some cases false decisions made by the user could not be reverted without starting the whole segmentation over. Even though the segmentation is done for one time step, all other time steps can be viewed as well or played in a video sequence. This time information has proven very helpful for the user when placing important landmarks, since they are not always visible or certainly visible in all time steps. As soon as the segmentation is done the user is supplied with a list of model based measurements that can quantify MV pathologies.

Part III

Methods and Evaluation

5 Objectives

After discussing the medical and technical background, this section will derive the final concept for this work. Ionasec [21] has shown that it is possible to segment the valve apparatus fully automatically. However, their method combines different fields like machine learning, tracking and shape models and is therefore very complex. Trying to reimplement or extend the method for evaluation in this master thesis might fail due to the necessary workload. Furthermore Inoasec's method is fully automatic but less accurate compared to the others. Also, in this thesis the annulus shape can be already obtained by the semi-automatic segmentation approach by Graser [6], which is integrated in the MITK Mitralyzer Plugin, and may be useful to automate the user initialization proposed by Pouch [5] and Schneider [22].

As stated in section 3.2.1, a large and variable collection of data is required to build a useful SP. After a first investigation, the data sets available for this work were found to be too few and the MV anatomy to be too variable. An SP created from the available data would likely not have been able to cover the variations of MVs that were not integrated in the SP. Consequently, an SP will not be used in this work.

An energy minimization will be utilized to optimize an initial segmentation. The technique is inspired by the *snakes* introduced in section 3.2.1. In this adaptation, the internal force is not based on an SP, it is more used as an urge of the model to extend itself into a specific direction. The internal and external force will be defined in section 6.3.8 so that they are in equilibrium, as soon as the model representation expresses a good delineation of the MV.

The segmentation approach in this work will be designed to enable a fully automatic segmentation. Nevertheless there should still be the option for a user to have influence on different steps of the segmentation process. From this the necessity arises, that the manual segmentation, as well as the semi-automatic segmentation, work on the same model representation and that the processing pipeline is broken down into single parts that can potentially be modified by a user.

The segmentation results shall then be quantified and compared to other methods. Also, the results should be usable in the assessment of a patient's MV.

The listing summarizes the objectives:

1. Implementation of an automatic method to segment the anterior and posterior leaflets of the MV, to capture patient-specific detail and establish correspondences on valves of different subjects.
2. Implementation of a manual segmentation method. This is necessary to build up a ground truth that can be used to measure the quality of the automatic process.
3. Combine the manual and automatic segmentation to a semi-automatic method.
4. The required user interaction to initialize or guide the segmentation should be minimized.
5. The accuracy of the segmentation result should be in the range of the variability of manual expert segmentations and other methods.
6. Clinically relevant measurements should be possible, based on the MV model.

5.1 Benefits from a Segmented Valve

The following will give another motivation from the surgeon's and researcher's point-of-view. The surgeons can benefit from a segmented MV model of a patient during intervention planning and surgery. TEE images (see section 3.1) can be used by the surgeon to determine MI severity and to plan necessary actions to reconstruct the valve. Generally, it is a challenging task for the cardiologist to find the correct position and angles of the US probe, to take significant 2D images of the valve that can be used for diagnosis. Also, it requires a lot of experience to understand the spatial relation between different 2D images.

Since the TEE imaging is capable of delivering 3D volumes, a 3D segmentation of the valve can supply the surgeon with numerous additional information. As shown in [6] the annulus can be segmented to find the best fitting annuloplasty ring prosthesis, obviating the need for intra-operative measurements. Traditionally, the surgeon measures the annulus intra-operatively at the open heart. Other applications could be to identify anomalies in the leaflet anatomy to determine which leaflet segment has to be treated for optimal results. Further, the closing capabilities and the motion sequence of the mitral valve over a cardiac cycle could be analysed. When replacing the chordae, the optimal chordae length could be measured to make the leaflets close properly. Wrong chordae length may lead to unnecessary strain at the chordae and leaflets that may result in rupture of the tissue. Additional measurements, like the extent of a regurgitation orifice

or leaflet thickness, may give more information about the severity of a patient's disease.

This is only a subset of possible measurements that could be done on a MV model. Datasets of MV models may give physicians the opportunity to do research in new directions and expand their knowledge about the MV anatomy and behaviour. The great obstacle at present time is that there is no ready solution to efficiently create MV models. It is possible to use common manual segmentation tools but this is very time consuming and a precise result is not guaranteed. A semi-automatic or even automatic process with satisfying accuracy could be utilized to create large collections of data in reasonable time. An open source solution of this process would make the results available to any researcher.

Probably, different experts will model the same valve differently. A machine based method would always create models with the same bias, if any. Furthermore, valve models could be used for simulation to predict the outcome of a surgery, e.g. by simulating how well a valve will work after a certain portion of a leaflet has been resected.

Time plays another important role during surgery and clinical routine. The structure of the MV is very variable and the segmentation of the annulus over one cardiac cycle has taken $551 (\pm 161)$ seconds for 25 time steps in [6]. For the whole MV it would be even more. This is inapplicable for the use during interventions. Thus, for intra-operative use, segmentation systems should be at least semi-automatic to save working time.

5.2 Manual Segmentation

Like other publications, this work will compare manual segmentations of the MV to the semi-automatic results. To do so within the Mitralyzer, the plug-in has to be extended to allow for leaflet segmentation. Working experimentally with QLab has shown that it is crucial to present the correct slices to the user that show the anatomic features of interest. Guiding the user through the volume is a primary factor in decreasing the time needed for a segmentation. If the guidance by the program is too weak, the user has to put too much time and effort in finding the slices independently. If the guidance is too strong and the presented slices do not contain the desired information, it is hard for the user to perform a good segmentation. The implementation of the manual segmentation will be described in section 6.1, following the objectives:

1. A previously segmented annulus model from the same TEE is required as input.
2. Split the segmentation process into separate consecutive steps.

3. Guide the user to easily find anatomically relevant slices of the volume.
4. Allow definition of leaflet tips.
5. Allow definition of the leaflet curvature.
6. Allow definition of the papillary muscle location.
7. Allow segmentation over a whole cardiac cycle (all time steps of a data set).
8. Allow that the user can go back and forth in the defined workflow at any time.

5.3 Semi-automatic segmentation

The approach in this thesis will be the partly reimplementation and refinement of Schneider's method [22]. It will be combined with an important initialization step that is found in Pouch's work [5], which is the initial separation of the leaflets. In their work this had to be done manually. Another important ingredient is the clustering of voxels to label them as blood pool or tissue. This was done for the whole image in Burlina's work [30] and a variation will be used in this work.

The objectives of the method described in section 6.3 are defined as:

1. Use a mitral annulus model as input.
2. Automate the leaflet separation (as manually done by Pouch).
3. Pre-process the image data for model initialization
4. Fit a 3D model into the image data to obtain a fine surface mesh of the valve.
5. A user should be able to guide the automatic modelling process.

6 The Segmentation Pipeline

This section will describe the steps of the segmentation pipeline. Section 6.1 will introduce the manual segmentation, including the structure of the 3D model representation. An overview of all the automatic pipeline's components can be found in figure 34. Section 6.2 shows the assembly of all components to a segmentation pipeline. Each component will be explained separately in section 6.3, focussing on what each component basically does.

Glossary The following sections will use a couple of terms that have more or less medical relevance and should be defined first.

- **Saddle Horn** - Peak of the anterior part of the annulus, which is close to the aortic valve.
- **Annulus Commissure** - Two points on the annulus ring that are closest to the actual commissure on either antero-lateral and postero-medial side.
- **Commissure** - The actual two commissure regions on antero-lateral and postero-medial side, where the leaflets join.
- **Coaptation Point** - Several points at each leaflet tip, that can touch the opposite coaptation point during valve closure.
- **Leaflet Point** - The leaflets' edges are defined by the annulus and coaptation points. The surface itself is defined by leaflet points. These points represent the leaflets as a spline on 2D slices.
- **Commissure Line** - The connection of a commissure point to its affiliated annulus commissure point.
- **Coaptation Line** A connection between both commissure points that runs through the center of the coaptation points.

6.1 Manual Segmentation

In order to create a manual segmentation of the whole valve, the Mitralyzer plug-in had to be augmented with the ability to segment the whole MV including the annulus, both leaflets and the papillary muscles. The design had to be made with respect to user-friendliness since it is desirable to support the user with sufficient usability to ease segmentation effort. As found in section 4.3 it is very important to supply the user with the appropriate slices of the volume to make a good segmentation possible.

The work-flow now consists of the following steps:

1. Annulus segmentations as described in chapter 4.1
2. Placement of commissure and coaptation points
3. Modelling of the leaflet shape
4. Placing papillary muscles

To allow modelling of the leaflets, the annulus model representation proposed by [6, 19] had to be extended. Figure 27 shows a schematic of the

new model. Each coaptation point is connected to one annulus point, commissure points are connected to three annulus points. The connection is subdivided to allow for modelling of the leaflets. The green points represent the commissure. From an anatomical point of view, the commissure on antero-lateral and postero-medial side is a region where both leaflets converge. It can hardly be represented by a single point, but for the sake of usability the user only has to define a point that lies within this region, rather than to define a small 3D region.

There are four commissure points, two on the annulus curve and two on the leaflet coaptation curve. The points that represent the anatomical location of the commissure are the ones on the coaptation curve. The other points should be set by the user on the annulus, but as close as possible to the actual commissure. The saddle horn should be set on the peak of the annulus curve, next to the aortic valve.

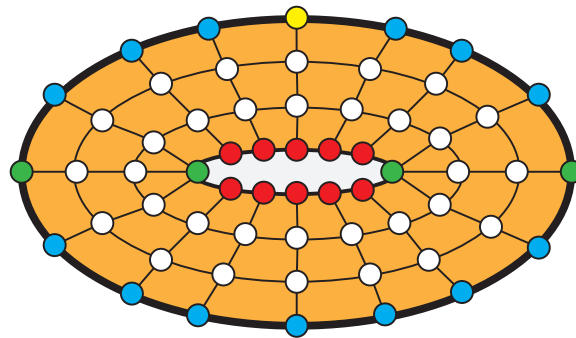


Figure 27: Schematic of the new model representing the annulus (blue), coaptation zone (red), commissure (green), saddle horn (yellow) and leaflets. Note that the schematic shows only two subdivision points for the leaflets (white) whereas the real model has five.




Figure 28: Schematic of the fan-like slicing.

Placement of commissure and leaflet coaptation points is done to define the structure of the coaptation or leaflet ends, respectively. The points are

placed in slices that are arranged in a fan-like manner (see figure 28). The outer slices are defined by the commissure points and the remaining slices are distributed in equal angles between them. All slices are fixed to the saddle horn. This slicing method is called *Fan-Slicing*. Except for the outer slices, an endpoint for the anterior and posterior leaflet can be set each, eventually forming the shape of an opened or closed MV. Note that one DOF gets lost for setting the support points since each point is bound to a fixed slice. The commissure points are initialized at a fixed location that does not necessarily fit the anatomy of the MV. The user can then rotate the outer slices around the fixed annulus point until the commissure is found on each side.

At this point, the model may have an unsatisfactory commissure line. If that is the case the user has the additional option to move the annulus commissure points along the curve to achieve a commissure line that smoothly extends the coaptation line (see figure 29).




Figure 29: The user can adjust the annulus commissure to define the commissure line. Left image: commissure before adjustment. Right image: commissure after adjustment.

The illustration shows a bad commissure line. After relocating the annulus commissure points, the coaptation line is smoothly extended by the commissure lines. The other annulus points are also moved to maintain an equal distribution along the annulus curve.

The placement for a commissure point and two coaptation points is illustrated in figure 30. The leaflets share the commissure points, so only one point can be set here in one slice. When placing the commissure points, the respective plane can be rotated. So, if the 2D slice for the commissure point looks like the one in the bottom row of figure 30, the slice can be rotated towards or away from the annular commissure until the leaflets touch and where a commissure point can be set reasonably. In the case of the coaptation points the slices are fixed and one point per leaflet can be set to allow for the modelling of both, an opened or closed valve.




Figure 30: Placement of a commissure point (top) and two coaptation points (bottom). The dashed line on the left represents the slice through the volume, resulting in a 2D view. The 2D view on the right shows tissue (gray) and blood pool (blank).




Figure 31: Leaflets are segmented separately. Note that, while the schematic shows two equal slices, the slices for two corresponding coaptation points can actually be different. The left slice (posterior) is taken from the green dashed line, the middle slice (anterior) is taken from the red dashed line.

Refinement of leaflet shape is done using the previous steps. The annulus and coaptation points are paired clockwise. Each annulus point is paired with a coaptation point that is directly connected to it via leaflet points. The image volume is then sliced so that each slice contains one of the control point pairs. The user can now view the different slices and form the leaflet shape by shifting control points to the leaflet curve. The coaptation points can also be refined in this step, even though they should already be in a correct position. A segmentation for an anterior and posterior leaflet

slice can be seen in figure 31.

On the right hand side of figure 31, the second anterior annulus point (excluding the annulus commissure) is connected to the second anterior coaptation point (including the commissure) and so on. The same is done for the posterior leaflet, resulting in two slices that are shown as a dashed red and green line. The start of the spline in each slice is fixed by the corresponding annulus point (white). Because the volume is sliced separately for the anterior and posterior part of the valve, this slicing method is termed *Double Slicing*. The reason why the fan slicing is not used here can be seen in figure 32. How much of the MV area is covered by the leaflet slicing depends highly on where the commissure points were set.




Figure 32: The fan slicing covers only a portion of the MV, depending on the commissure points' position. The covered area is displayed in orange.

Placing papillary muscles is the last step. Two slices help the user to define the centers of the papillary muscles. One slice is aligned to the commissure line and the other orthogonally to the first one. The correct location is hardly visible in US images in most cases, so the segmentation should just be a rough guess. Figure 33 shows the concept, where the user has to define two points that approximate the location of the papillary muscles.



Figure 33: The bulges at the bottom of the ventricle can be identified as the papillary muscles. Red dots: Points estimating the papillary muscle location.

The user has to repeat these steps for a sufficient number of time steps until the linear interpolation between the time steps delivers a satisfying result.

6.2 Pipeline Assembly

This section will describe how the final segmentation pipeline is assembled, each component is described in section 6.3. The flow chart in figure 34 gives a quick overview of the full pipeline. There are four paths that describe different pipelines.

- A: A leaflet separation step is used to refine the clustered 3D graph cut by separately masking the leaflets and applying the graph cut to each masked region. This is called the *masked pipeline*.
- B: Skipping the masking step and using the initial clustered 3D graph cut from step four in the pipeline for further processing, results in the *unmasked pipeline*.
- C: The model fitting is done completely automatically, which is termed the *auto* configuration.
- D: The model fitting is guided by the user. The user can either set the commissure points or the commissure and coaptation points. These configurations are called *semi-1* (commissure points defined) and *semi-2* (commissure and coaptation points defined).

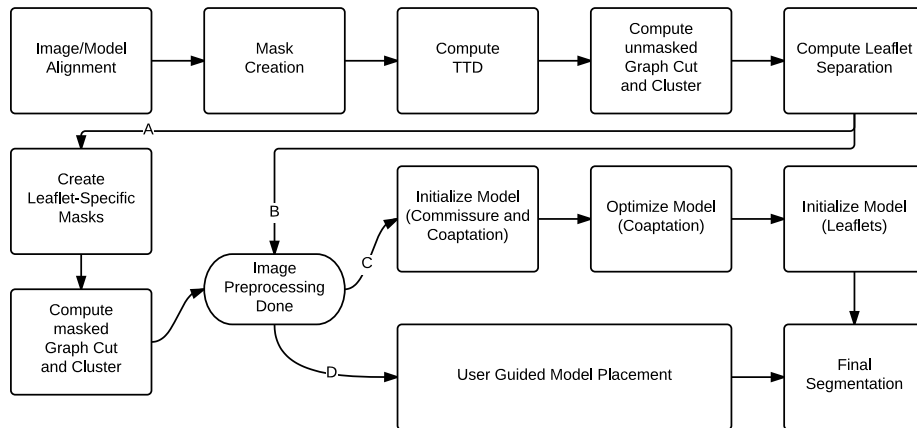


Figure 34: Segmentation Pipeline, showing different configurations (A, B, C, D) of the (semi-)automatic pipeline.

A 3D+t US image of the MV region and an annulus segmentation of this image are needed as input. The annulus model and the image are aligned as described in section 6.3.1. Next, a mask is created that restricts image

processing to a minimal ROI (see section 6.3.2). The TTD is then applied to the masked region of the input image (see section 6.3.3). The TTD results are utilized to create a *clustered 3D graph cut* image (section 6.3.5). A MIP of that clustered 3D graph cut is then used to do a leaflet separation with a *2D graph cut* (section 6.3.6). The leaflet separation will be needed in the model initialization process.

At this point, the image preprocessing is finished, when using path B in the flow chart, which would be the *unmasked* result with a 3D graph cut that is not based on leaflet separation. Following path A, the separated leaflets are used to divide the initial ROI into two leaflet specific ROIs (optional step of section 6.3.6). A clustered 3D graph cut is then applied to both masked regions individually (optional step of section 6.3.5). The result of these two clustered 3D graph cuts is merged to a single image. Thus, the masked pipeline requires extra computational effort, since two additional clustered 3D graph cuts have to be processed. This would be the *masked* result. Depending on path A or B the masked or unmasked clustered 3D graph cut is used during the subsequent steps. Table 2 Summarizes the input and output of each preprocessing component.

	Image/Model Alignment	Mask Creation	Compute TTD	unmasked clustered Graph Cut	Leaflet Separation	Leaflet Mask Creation	masked clustered Graph Cut
	<i>general masked and unmasked pipeline steps</i>					<i>additional masked pipeline steps</i>	
Input 1	original US image	aligned US image	aligned US image	TTD Image	MIP of clustered unmasked 3D Graph Cut	2D Graph Cut (Leaflet Separation)	3D Mask anterior/posterior
Input 2	original Annulus	aligned Annulus	3D Mask	3D Mask	MIP of 3D Mask	3D Mask	TTD Image
Input 3				aligned Annulus	aligned Annulus		aligned Annulus
Output 1	aligned US image	3D Mask	TTD Image	unmasked 3D Graph Cut	2D Graph Cut (Leaflet Separation)	3D Mask anterior/posterior	masked 3D Graph Cut (merged anterior/posterior)
Output 2	aligned Annulus			clustered unmasked 3D Graph Cut			clustered masked 3D Graph Cut (merged anterior/posterior)

Table 2: Summary of each preprocessing component's input and output.

Image preprocessing is now done and the model has to be placed, based on that preprocessing. There are also two paths available, a fully-automatic process (C) and a user guided process (D).

Path C initializes the commissure and coaptation points of the model (section 6.3.7), based on the masked/unmasked clustered 3D graph cut and

its MIP. After that, the model is optimized as in section 6.3.8. For this, the TDD image is used. The optimization works for coaptation points and leaflet points based on the fan slicing. But, the fan slicing is not used in the manual leaflet segmentation. This is why, in a last step, the leaflets have to be initialized based on the double slicing method which was explained in section 6.1. The step is necessary to bring the model into a structure that is editable in the manual segmentation process. A quick look at intermediate results of the pipeline is given in figure 35.




Figure 35: Intermediate results of the pipeline. TEE input image (1). TTD (2). 3D graph cut (3). Clustered 3D graph cut (4). Initialized model (5). Optimized model (6). Final model and TEE input image (7).

When choosing the user guided segmentation (D), the same steps are done, but the user can check and edit each step. The process is illustrated in figure 36.

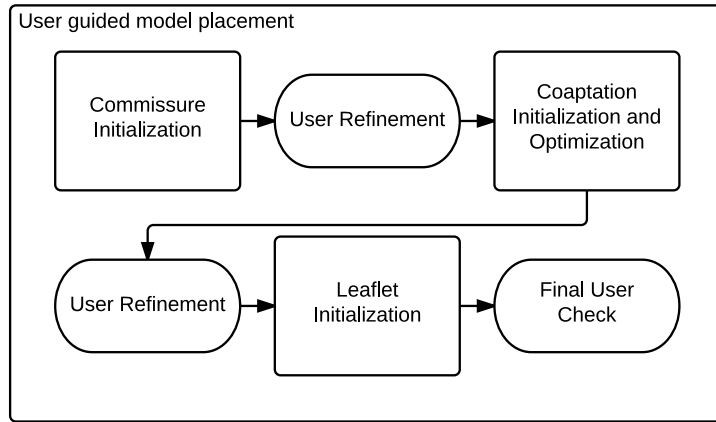


Figure 36: Segmentation Pipeline.

The commissure points are placed automatically, then the user can adjust both points. Based on the user input, a set of coaptation points is proposed to the user. Again, the points can be edited. Finally, the leaflet points are initialized to undergo a last user check.

This semi-automatic segmentation is designed to achieve segmentations with the same quality as a manual segmentation, but with less user effort, since checking the model points for being in a valid location consumes less time than setting each single point manually. Obviously this advantage relies on the initialization quality of the algorithm. If every point has to be relocated, no time is saved.

6.3 Semi-Automatic Pipeline Components

This section provides a detailed description of all components that are used in the segmentation pipeline. The order of components here already gives a rough idea of the final pipeline. The complete pipeline will be described in section 6.2.

6.3.1 Image and Model Alignment

The first step of the pipeline is an alignment of the image and model. A plane fitted into the annulus shape is regarded as the *annulus plane*. The normal of that plane is the *annulus normal*. Later steps of the pipeline traverse voxels in the direction of the annulus normal. For reasons of performance it is beneficial to align one axis of the US input image with the annulus normal. In this work the model was aligned to the image z-axis, or axial direction respectively. Now, traversing voxels in the direction of the annulus normal is the same as traversing along the z-axis and no interpolation in 3D space is required.

6.3.2 Masking the Volume

In order to restrict the time consuming image processing to a small volume, the image volume is masked based on the annulus shape. This ROI should contain possible locations of the leaflets. The annulus plane is extruded in positive and negative axial direction by a size of half of the annulus diameter. The 3D ROI is shown on the right of figure 37. The center image shows an axial view, here the mask fits perfectly into the annulus shape. The image on the left shows a sagittal view which clarifies the extrusion. All further processing is restricted to this ROI.




Figure 37: Annulus ROI displayed as highlighted region and volume. Left: Sagittal view. Middle: Axial view. Right: 3D volume with annulus curve.

6.3.3 Thin Tissue Detector

The *Thin Tissue Detector* (TTD) proposed by Schneider et al. is designed to find locations of thin tissue in ultrasound images. The following description of the algorithm is based on [10]. The TTD works on a 3D volume Ω and uses basic, low level and local image processing techniques. The final TTD term has three components ψ , θ , β and two parameters σ and p . ψ is obtained by computing a Gaussian smoothing with variance σ on the gradient magnitude image of Ω :

$$\psi = G(\sigma) * |\nabla\Omega|$$

where σ is set to half of the expected tissue thickness in voxels and the kernel size to $4\sigma + 1$. For θ and β the image is split into a set of cubic neighbourhoods N of edge length p . N_i consists of all voxels n_j within the cube centred at v_i where $j \in [1, J]$ and $J = p^3$. θ_{v_i} is then the mean angle between all gradient vectors in N_i on the smoothed image $S = G(\sigma) * \Omega$:

$$\theta_{v_i} = \left(\sum_{j=1}^{J-1} j \right)^{-1} \left(\sum_{a=1}^{J-1} \sum_{b=a+1}^J \cos^{-1} \left(\frac{\nabla S_{n_a} \cdot \nabla S_{n_b}}{|\nabla S_{n_a}| |\nabla S_{n_b}|} \right) \right)$$

β sums up the net flow of a neighbourhood N_i to its neighbours with the gradient vector of each voxel as flow direction. Each boundary of N_i has an inward directed unit normal denoted as $\mathbf{f}_{q,norm}$ with $q \in [1, 6]$ for all sides of a cube. Voxels adjacent to a face q are denoted f_k with $k \in [1, K]$ and $K = p^2$.

$$\beta_{v_i} = \sum_{q=1}^6 \sum_{k=1}^K \left(\frac{\nabla S_{f_k}}{|\nabla S_{f_k}|} \cdot \mathbf{f}_{q,norm} \right)$$

The final TTD value is then computed as a product of θ , β and ψ with β and the result is normalized to $[0, 1]$:

$$TTD_{\Omega} = \theta \cdot \beta \cdot \psi$$

Higher TTD values indicate that a voxel is more likely at a thin tissue location. Figure 38 shows the components and results of a thin tissue detection.




Figure 38: Component value images of the Thin Tissue Detector. input image (1), ψ component (2), θ component (3), β component (4), result (5).

In this work $\sigma = 2$ and $N = 3$ was used, just like in [10].

6.3.4 Maximum Intensity Projection

As the TTD represents thin tissue voxels with a high intensity, the Maximum Intensity Projection (MIP) is a simple way to transfer the 3D representation of an image to a 2D representation. All voxels are projected on a plane in axial direction. Each target pixel contains the maximum intensity value of all voxels that were projected on it (see figure 39 for an illustration).




Figure 39: Maximum Intensity Projection from 2D to 1D.

It seems that a lot of information is lost with a MIP, but as the focus lies on high intensity locations and the annulus normal is used as the projection direction, the result can yield important information about the valve state and shape. Figure 40 shows a MIP of a TTD image. The centered black region can be interpreted as either very thick tissue or the absence of tissue. Another option is that the leaflets coapt in this region and form a thicker region of tissue, but this usually results in a dark gray area instead of mere black. Since it is anatomically not plausible to find thick tissue at this location, it can be concluded that the dark area shows an orifice.




Figure 40: MIP of a TTD image.

The observation that the MIP gives a simplified, yet helpful view on the valve will be exploited at a later point of the pipeline.

6.3.5 Clustered Graph Cut

The graph cut is used to find a surface that contains the mitral leaflets, based on the TTD image. The special search space proposed by Schneider (see section 4.2.2) and an existing implementation of Kolmogorov’s graph cut algorithm [17] is used here. The node values are taken from the TTD image only, not the drive image suggested by Schneider, and are connected to sink an source as proposed by Schneider in [22]. The parameters used are shown in the following table. The values differ from the ones proposed in [10], where the graph cut was used to find the leaflets during systole only.

Parameter	ω	ω axial	α	α axial	Θ_p	Θ_s	R
Value	1	25	25	25	256	128	64

Table 3: Refer to [10, 22] for the meaning and usage of these parameters. The first four parameters make the graph cut more flexible in axial direction and less flexible in the axial plane. The last three parameters mean that the Schneider search space has a resolution of $256 \times 128 \times 64$ voxels.

While Schneider uses an active contour to separate the leaflet tissue from the blood pool in the resulting 3D graph cut surface, a clustering is done here that yields a *clustered 3D graph cut*.

The initial 3D graph cut surface is transformed into a binary image, where the surface is represented as foreground with a thickness of one voxel. A morphological dilation is applied on the foreground. The radius of the dilation has to be chosen greater than half of the expected leaflet thickness. This makes sure that the dilated foreground includes voxels that represent the blood pool, even if the graph cut lies completely within tissue. A *k-means* clustering is then applied on all foreground voxels, labelling three classes. The classes are to be read as *tissue*, *uncertain* and *blood pool*. In figure 41 the original and clustered 3D graph cut are compared. The green area is the union of the tissue and "uncertain" regions.

The original graph cut reaches into the blood pool. The clustering is used to label these undesired foreground voxels as blood pool and eliminate them. If there were no foreground voxels of the original graph cut within the blood pool, the clustering would find false negatives (falsely labelled as blood pool), since no class remains empty after the clustering. Because of that, the dilation is done. The uncertain region is necessary to avoid a hard transition between blood pool and tissue. For the optimization process, that will be introduced in section 6.3.8, it is beneficial to know where tissue is located with high certainty. The boundary between tissue and blood pool will then be detected by the optimization itself.




Figure 41: Original (red) and clustered (green 2D, gray 3D) graph cut. Each column showing a 2D slice and a full 3D surface.

As visible in the right column, there are still artifacts left (see the right leaflet which has a gap in it).

Another issue with the graph cut is, that the cut surface may run through the blood pool instead through tissue, if this path is significantly shorter or cheaper in terms of node capacity. This may happen if the leaflets are in open valve state and if the search space overlaps itself, meaning that the arcs defining the space include parts of the leaflet on the opposite side (refer to figure 19). Figure 42 shows the effect (red lines are actual segmentations done by the algorithm). In the left image the graph cut connects both leaflets instead of finding their respective tip. The radius cannot just be decreased as the center image shows. That would lead to the case that parts of the leaflets are not covered by the search space. The desired segmentation was manually drawn in the right image. Since the leaflet length is not available prior segmentation and also varies within different leaflet segments and time steps, the radius would have to be adjusted manually to best fit the data at hand.




Figure 42: Radius too large (left). Radius too short (middle). Desired segmentation (right).

To overcome this problem, the 3D graph cut could be applied to separate regions that only contain one leaflet. An automatic separation of the leaflets would be required for this. An approach to achieve this is specified in the following section. If no leaflet-specific masks are used for the graph cut, it will be termed *unmasked 3D graph cut*, if leaflet-specific ROIs are used to refine the graph cut, it will be called *masked 3D graph cut*.

Using the masked 3D graph cuts is *optional* and required in the second additional step of the masked pipeline (path A in figure 34).

Figure 43 shows the results for a masked and unmasked 3D graph cut. The red line is the original result and the leaflets are connected without reaching their respective leaflet tip. The green line shows the result with a leaflet-specific ROI and the tip of the left leaflet has been found successfully.




Figure 43: Comparison of masked (green) and unmasked (red) graph cut. Overlapping results (dark green). Masked region (yellow).

6.3.6 Leaflet Separation

Separating the valve in two leaflet parts also uses a graph cut. In this case the cut is done on a 2D MIP of the input image, which can be a TTD image or clustered graph cut image, resulting in a *2D graph cut* and designed as follows:

1. A **Region of Interest** defines the space where the leaflets can be found in all time steps.
2. A **MIP** is applied to the defined ROI on the input image in axial direction.
3. A **Graph Cut** is applied on a 2D graph constructed with the values given by the MIP. Nodes at the edge of the ROI are connected to either source or sink, depending on the annulus segmentation. The cut line roughly delineates a coaptation line.
4. **Create a Mask** from the initial annulus ROI. The ROI is cut in two by the cut line. (This step is *optional* and required in the first additional step of the masked pipeline (path A of figure 34))

Figure 44 shows a result of this graph cut. The annulus plane is cut in two regions as an estimate for the posterior (highlighted region) and anterior leaflet. If desired, the 2D regions are then extruded in the manner of the initial ROI to create a 3D ROI for each leaflet. The borderline between these two 2D regions or the border-surface between the two 3D regions is termed *leaflet separation*.




Figure 44: Leaflet ROIs displayed as region and volume. Left: Cut regions of MIP. Right: 3D ROI.

The topology of the graph is illustrated in figure 45. The yellow dots are the commissure points that were set during the annulus segmentation. Nodes inside annulus area are connected as a common grid in a 4-neighbourhood, the grid is aligned to the input image axes. Voxels at the edge are connected to either sink or source. The red line covers voxels that are on the saddle horn side of the commissure line, they are connected to source. The remaining voxels (green) are connected to sink.




Figure 45: 2D Graph topology.

The cut is determined by the distribution of sink and source nodes, since the cut is drawn to the boundary between sink and source nodes. This means that the cut line will always run through the annular commissure. The edge capacities for the graph can either be taken from the original image, TTD values or clustered graph cut values, leading to different results.

Figure 46 compares results of a TTD and clustered graph cut input image.

The yellow dots show where commissure points were set during the annulus segmentation. These locations separate the edge pixels into source and sink nodes.




Figure 46: Leaflet separation. Based on clustered graph cut (green). Based on TTD (red).

Using the clustered input image can better ensure that the cut runs through the blood pool, as the intensity difference between two labelled regions is higher than in the smooth gradient of the original TTD image, where low intensity regions may be skipped. In some cases, the mere clustered image is not enough to provide a satisfying result. Knowledge about the MV anatomy can be used to further optimize the graph cut. Carpentier [8] delineates the coaptation with a shape similar to the posterior part of the annulus, with a distance of 5-6 mm to the annulus curve. Using this information, the node capacities of the graph can be adjusted. Figure 47 shows an illustration of the weights that are applied.




Figure 47: Leaflet separation curve weight illustration (left). Separation with curve weight (green) and without curve weight (red.)

Darker regions in the mask mean, that the node capacity at this location is reduced more than lighter regions. White regions are not modified at all. This way the graph cut is drawn towards the anatomically more reasonable location. The capacity modification is done by the following rule:

$$Capacity_i = \begin{cases} Capacity_i, & \text{if } \Delta_P(i) \geq 2 * \Delta_p \\ Capacity_i * \left(\frac{|\Delta_p - \Delta_P(i)|}{2 * \Delta_p} + 0.5 \right), & \text{otherwise} \end{cases}$$

where $Capacity_i$ is the capacity of a graph node i , $\Delta_P(i)$ is the distance of a node or pixel i to the posterior annulus and Δ_p is the expected distance of the coaptation line to the annulus. The penalty factor is mapped to $[0.5..1]$ so that the node capacity is at most halved.

Due to the advantages described above, the final pipeline will use a MIP of a clustered 3D graph cut and the modified node capacities to find the leaflet separation.

6.3.7 Model Initialization

The model initialization is an important step for the final outcome of the model optimization, since the optimization relies on the assumption that the model lies within areas of high TTD intensities. The MV model presented in section 6.1 has to be fit into the preprocessed volume data. The initialization is done in three steps:

1. Find commissure points
2. Find coaptation points
3. Find leaflet points

Step one and two will be initially performed on a MIP of a clustered 3D graph cut. These steps are both followed by a back-projection into 3D space, so that the points are correctly defined in the original volume.

Commissure Initialization

The commissure points are searched along the line that separates the two leaflets. Beginning at the annulus commissure the classes of the clustered image are checked until a transition from tissue to blood pool is detected. If no such transition was found until up to 10 mm away from the annulus commissure, the location of a detected transition with lower priority is used. If no transitions were found, the last checked point on the cut line is chosen as a commissure point. Since there are three classes and all possible transitions of descending class order have to be considered, transitions are prioritized in the following way:

1. Tissue - Blood Pool

2. Tissue - "uncertain" - Blood Pool
3. "uncertain" - Blood Pool
4. Tissue - "uncertain"

If a transition was found, the location of the higher intensity class is chosen as the commissure point. Figure 48 illustrates this step. The commissure points are now located on the MIP plane of the clustered 3D graph cut image. To find their final position in 3D space, they have to be projected back to the original clustered 3D graph cut image.

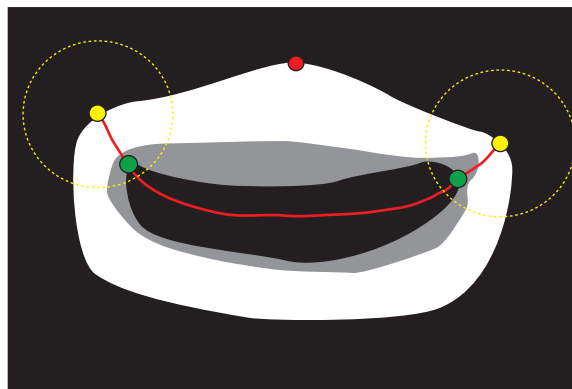


Figure 48: Commissure initialization schematic. Yellow dots: Annulus commissure. Green dots: Commissure. Red line: leaflet separation (2D graph cut). Red dot: Saddle horn.




Figure 49: Coaptation initialization schematic, step 1. Blue dots: Coaptation points. Yellow dashed line: Fan slices.

Coaptation Initialization

Initialization of the coaptation points depends on the commissure locations. The angle formed by saddle horn and commissure points is divided into sections with equal angles at the saddle horn. Each point where the

leaflet separation and these rays intersect is used as a first initialization for a coaptation point of each leaflet.

As shown in figure 49, the space between the two commissure points is equally subdivided. The coaptation points of each leaflet are positioned with a small offset to their neighbours. In a second step, the coaptation points move on their respective slices towards their leaflet root on the annulus. Again, transitions between the three classes are considered to decide where the movement stops. Since the search direction is inverted, the priorities are also inverted:

1. Blood Pool - Tissue
2. Blood Pool - "uncertain" - Tissue
3. Blood Pool - "uncertain"
4. "uncertain" - Tissue

The final location is taken from the higher intensity class location. If no transition is found, the coaptation point remains in its position of the previous step. The resulting points depicted in figure 50 are also back-projected into 3D space to settle the final location.




Figure 50: Coaptation initialization schematic, step 2. Blue dots: Coaptation points. Yellow dashed line: Fan slice. Blue line: Orifice outline.

Leaflet Initialization

As a last step, the connection from each coaptation point to its root is subdivided by leaflet points. These points are also projected into their final location by looking at the 3D graph cut surface and the clustered 3D graph cut volume. The leaflet points are at first distributed over the direct connection from the coaptation point to the root, as illustrated in figure 51. For each leaflet, voxels on an arc (similar to Schneider's search space) are checked for the presence of the 3D graph cut surface and the class of the

clustered graph cut. The voxels are sorted into the following categories depending on what information is found at their location, with descending priority:

- Tissue and graph cut
- "uncertain" and graph cut
- Tissue
- "uncertain"

The location of all voxels in the first group is averaged and set as the new location for the respective leaflet point. If the first group is empty, the second group is used and so on. If all groups are empty, the leaflet point remains at its initial position.

All leaflet points that belong to one coaptation point are then checked for validity. The angle between the vectors of each leaflet point to its neighbours has to be above a threshold n and the angle between the vectors of each leaflet point to the leaflet tip and annulus has to be above a threshold o . If one of the angles is below this threshold, the point is treated as an outlier. In this work $n = 90^\circ$ and $o = 100^\circ$ was used. All outliers are relocated by creating a spline, using all valid points. The outliers' locations are then taken from the interpolated spline.




Figure 51: Leaflet initialization schematic. Red solid line: Graph Cut. Blue dashed line: Connection of leaflet tip to root. Yellow dot: Point on the annulus curve. Blue dot: Coaptation point. Yellow dashed line: Search arcs. Red dots: Leaflet points.

The model is now initialized but one must not expect this to be a satisfying segmentation. The goal of this step is to put the important landmarks of the 3D model into reasonable locations, so that the following model optimization can work reliably.

6.3.8 Model Optimization

The initialization result shown in figure 51 is not optimal. Because of the MIP information about the leaflet tip location has been lost (consider a vertical projection). Also, the initialization has prioritized the tissue class of the clustered 3D graph cut. The leaflets might be too short now, because the actual transition from tissue to blood pool is somewhere within the uncertain class. An optimization process shall now use 3D information to find the tip of the leaflet. Additionally, 4D information will be used to prevent outliers. If the valve is in closed state, the algorithm should lead the leaflets to coapt.

So there are two main goals for the optimization: Finding the outermost leaflet tip during open valve state and find a coaptation of both leaflets during closed valve state. The process uses the intensity values of the TTD image, since the three classes of the clustered graph cut would be too restrictive. The movement of the model points is limited to their respective fan slices. This means that each coaptation point can coapt with its opposite point.

The method used is similar to an active contour method. For each coaptation point and the leaflet points belonging to it, a force is computed. This force moves the points in one iteration to a location where less force is expected in the successive iteration. Iteration stops when the summed force of all points falls below a threshold. In other words, the norm of the force $F(x)$ is minimized.

$$\arg \min_x |F(x)|$$

where

$$F(x) = (E(x) + U(x)) * s \quad (1)$$

where x is a vector containing all landmark points of the valve model (coaptation, leaflet and annulus points), $E(x)$ is an external force, $U(x)$ an internal force and s the step size. The minimization is done in an iterative process that follows the scheme presented in figure 52. A detailed presentation of the activities in this flow chart will follow after a rough overview.

After the model initialization, the model is checked for outlier points which are corrected if necessary. Next, forces for each point are computed. The forces yield a 3D vector for each point, that is then used to move the point to a location where a weaker force is expected in the next iteration. The convergence is checked by summing up the norm of the vectors of all points of all time steps. If the averaged sum is below a threshold, the iteration stops. Otherwise the next iteration starts. The forces are computed separately for each time step, only the outlier correction uses information of all time steps.

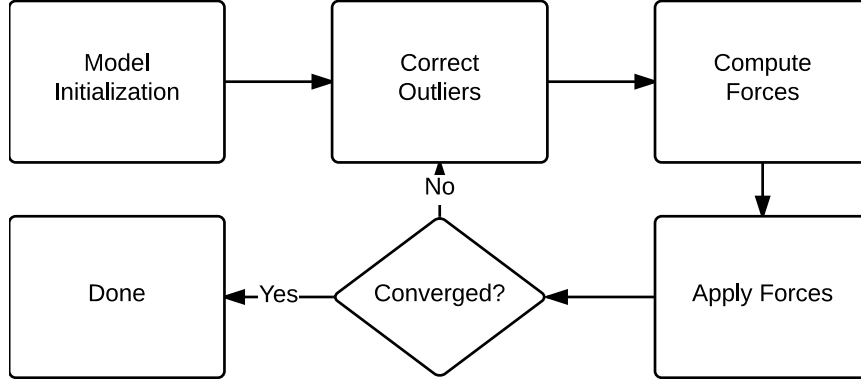


Figure 52: Flowchart of the optimization process.

The external force is based on the 2D gradient magnitude at each point's location in its affiliated fan slice.

$$E(x) = \nabla_{x,norm} * \frac{(r - (1 - |\nabla_x|))}{r} * (1 - \sqrt{I(x)}) * w_e \quad (2)$$

The gradient vector is initially normalized. The second term takes the gradient magnitude into account, which can be adjusted by $r \in [1..∞]$. If $r = 1$, the normalization of the vector is cancelled. For an $r > 1$ the influence of the gradient magnitude is lessened. The TTD values can vary in regions where tissue is detected and actually present, but the force should not be confused by those irregularities. The last but one term ensures that the external force has less power in regions of high intensity. This way gradients in tissue regions can be neglected, whereas gradients in darker regions, for example at the edge to the blood pool, have a strong influence and can keep the points inside the tissue region. w_e is a weight factor.

The internal force is defined as

$$U(x) = (P(x) + S(x) + G(x) + M(x)) * w_u \quad (3)$$

$P(x)$ is a *Push Force* that makes the leaflet stretch. The connected points p of a leaflet in one fan slice are enumerated from 0 to $n-1$, where 0 is the annulus point, $n-1$ the coaptation point and the remainder are the leaflet points. $P(p_0) = 0$, since the annulus point is fixed, for $i > 0$ the push force is defined as

$$P(p_i) = (p_i - p_{i-1})_{norm} * \frac{L_{mean}}{|p_i - p_{i-1}|} * w_p \quad (4)$$

with

$$L_{mean} = \sum_{j=1}^{n-1} |p_j - p_{j-1}|$$

The push direction is the normalized vector to a point from its preceding neighbour. The push strength is determined by the average segment length and the length of the current segment. End points of shorter segments get pushed harder than end points of longer segments. w_p is a weight factor.




Figure 53: Push force illustrated. Points get pushed away from a preceding neighbour, stretching the leaflet.

To preserve an equal point distribution, the *Smoothing Force* $S(x)$ moves points towards their farther neighbour. This force is only applied to the leaflet points so that $S(p_0) = S(p_{n-1}) = 0$. For $0 < i < n - 1$ let p_f be the neighbour of p_i that is further away and p_c be the one that is closer to p_i . Then

$$S(p_i) = (p_f - p_i)_{norm} * \frac{|p_f - p_i|}{|p_c - p_i|} * w_s \quad (5)$$

The normalized direction into which a point has to move is set towards the neighbour that is further away. The distance is based on the ratio of the distances to both neighbours. w_s is a weight factor.




Figure 54: Smooth force illustrated. Points move towards their farther neighbour, smoothing the curve and equalizing point distances.

Since the leaflet tips are usually directed towards the ventricle, a *Gravity Force* $G(x)$ is assigned to each leaflet tip, such that $G(p_i) = 0$ for $0 \leq i < n - 1$. Considering that the valve model and image are aligned in axial direction, the gravity has a fixed direction with weight w_g .

$$G(p_{n-1}) = \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix} * w_g \quad (6)$$

A last force is the *Momentum* $M(x)$ of a point. A portion of the total force $F(x)$ of the previous iteration is added to the current iteration to overcome small valleys in the TTD intensities. w_m determines the momentum's strength.

$$M(x) = F_{prev}(x) * w_m \quad (7)$$

$F(x)$ is now capable of optimizing the valve model in open valve state. To achieve a coaptation of both leaflets in closed valve state the forces for the coaptation have to be treated specially. Let p_i be a coaptation point and p_j be the opposite coaptation point in that fan slice. Then the adjustment is done for all coaptation points as follows

$$F(p_i) = \begin{cases} F(p_i), & \text{if } |p_i - p_j| > d \\ F(p_i) * (0.5 + a) + F(p_j) * (0.5 - a) + g, & \text{otherwise} \end{cases} \quad (8)$$

where

$$a = \frac{|p_i - p_j|}{2 * d} \text{ and } g = \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix} * w_{ga}$$

and d is a threshold that activates the coaptation force. The closer the coaptation points get to each other the more they share their force, adjusted by a . If the points end up in the same location they are finally treated as one point and a coaptation of both leaflets was formed. As soon as the points coapt, their behaviour changes and the gravity has to be adjusted, which is done with g . The reason for this is that the tissue region where the leaflets coapt is thicker. Consequently the TTD values are lower and the gravity has to be reduced, to prevent the coaptation points from dropping out of the thin tissue region.




Figure 55: Coaptation force illustrated. Coaptation points of both leaflets in the same fan slice move closer, driven by push and gravity forces. Finally, they coapt and represent a connection of the leaflets.

Another treatment that is applied only to coaptation points is the outlier correction. This process is meant for severe initialization errors or coaptation points that run into unreasonable locations during the optimization process. Before the forces are computed, the curvature among all coaptation points of a leaflet is checked within a time step. This is done by processing the angles of the vectors from each coaptation point to the coaptation points of the neighbouring fan slices. If this angle is below a threshold, the coaptation point is an outlier. The threshold in this work was set to 60° .

After checking all time steps, a parametric curve is built up for each coaptation point from all valid time steps of that point. The outlier points are then corrected by using the interpolated position at the corresponding time parameter. There is one parametric curve for x,y and z dimension each. Figure 56 shows an example with eight time steps. The point was found to be an outlier at time step two and six. The remaining points (now labelled from 1 to 6) are used to find a cubic curve that fits the available points. In order to simulate a time loop, the last time step is copied in front of the first one and vice versa. This ensures reasonable results for outliers at the first or last time step under the assumption that the first and last time step represent the start and end of a cardiac cycle, because otherwise no information would be available outside the parameter range.

Since the given points do not necessarily fit into a cubic curve, singular value decomposition (SVD) is used to find a best fit. To get the outliers to their original position in time, their parameters are now 1.5 and 4.5 as in figure 56 and the parametric curve returns their new location. To achieve a smoother movement of the coaptation points over time, the valid points are also relocated so that

$$p = (1 - f) * p_{valid} + f * p_{smooth}$$

where f is a factor that controls the influence of the smoothed point.




Figure 56: Outliers are fit into a parametric curve defined by valid points.

Recall that the outlier detection is done with information from one time step, while the correction takes all time steps of one point into account.

The convergence threshold was set to 0.03 and all other parameters of the optimization were defined as follows:

Parameter	s	w_e	r	w_u	w_p	w_s	w_g	w_m	w_{ga}	d	f
Value	0.03	1	10	1	0.2	1	2	0.1	-0.2	5	0.003

Table 4: Optimization parameters.

7 Implementation

This section will describe the implementation of the pipeline proposed in section 6 and other components developed for this work. Section 7.1 is about the manual leaflet segmentation that was added to the annulus segmentation process existing in the MITK Mitralyzer. In section 7.2 the structure of the main part of this work, the semi-automatic segmentation, will be explained and 7.3 describes the augmentation of another existing plug-in that is used to display measurements to the user.

7.1 Mitralyzer Extension

The Mitralyzer plug-in presented in section 4.1 is the result of the works of [6, 19]. It was augmented with the functionality to segment the mitral leaflets and papillary muscles as well, so that it could be used for ground truth generation for the evaluation (see section 6.1). Further, it is closely connected to the automatic part of the newly created SeMitralyzer plug-in and can thus be used for semi-automatic segmentation.

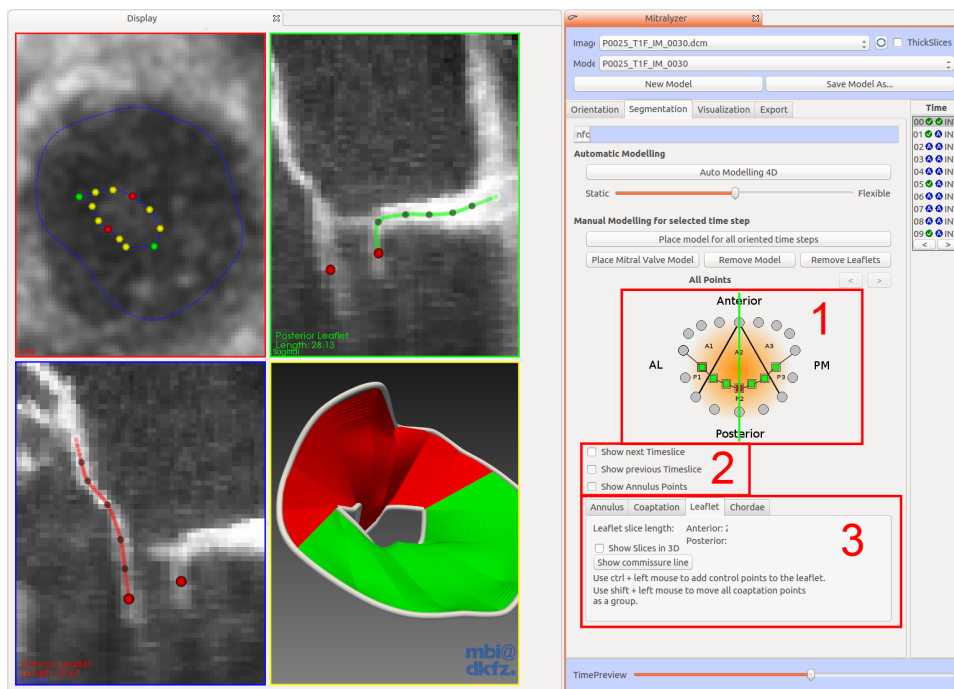


Figure 57: GUI of the extended Mitralyzer.

Figure 57 shows the new GUI of the Mitralyzer. The mitral model scheme has been updated (1) and now contains additional points that represent the commissure and coaptation points. It shows a closed valve but applies to diastole as well, because by selecting a coaptation point in the scheme

will select the affiliated coaptation point for each leaflet. Like this, an open valve can be modelled as well. In (2) the previous and next time step of the current points currently displayed can be shown to support the user. The tabs in (3) are used to switch between the different parts of the segmentation. If the *Annulus* tab is active, the user can model the annulus as in the old Mitralyzer version. The *Coaptation* tab allows modifying the coaptation and commissure points, the *Leaflet* tab is used to delineate the leaflet shape and the papillary muscles can be placed via the *Chordae* tab.

After an annulus ring was modelled with the given Mitralyzer functionalities, the user can now adjust the annulus commissure, to form a reasonable commissure line. Figure 58 shows how the anterior and posterior leaflet connection is determined by the annular commissures and by the leaflet commissures. The commissure line can be edited by moving this point.




Figure 58: The user can adjust the annulus commissure to define the commissure line.




Figure 59: Fan Slicing. Annulus scheme selection (left). 3D view displaying the model-slice correspondence (middle). Slice view to adjust the selected point (right). Commissure point selected (top). Coaptation point selected (bottom).

The placement of commissure and coaptation points can be seen in figure 59. The new Mitralyzer GUI shows an extended model scheme that can be used to select the points that represent the leaflet ends. In the top row, a commissure point is selected. Since both leaflets share one commissure point on either side, only one point has to be set here (green dot). The center image shows the selected slice in 3D, the 2D view is on the right. As described in section 6.1, the slice can be rotated about the saddle horn until a region that contains the commissure appears.

A coaptation point was selected in the bottom row. Each leaflet owns one coaptation point per slice, so the user has to set two points here (red for the anterior leaflet and green for the posterior leaflet).

When switching to the *Leaflet* tab the user will see two slices, created by the double slicing technique. Figure 60 shows both slices and the 3D view that clarifies the double slicing. Each leaflet is outlined in a separate slice, but the coaptation points (red dots) are present in both to give the user more orientation. These points can be edited in this view, but it is not recommended since the fan slicing is much more suitable to locate the coaptation. The small, dark dots represent the leaflet points. They are connected by a spline and can be relocated by the user to delineate the leaflet. The root of the spline is fixed to the annulus point that belongs to this slice.




Figure 60: Leaflets are segmented in separate windows. The 3D window shows the double slicing. The length of the spline is also displayed to support the user to maintain an anatomically reasonable leaflet shape.




Figure 61: Manual segmentation of a time step prior valve closure.

The Mitralyzer plug-in is also used during the semi-automatic segmentation process. It covers the manual part, while the SeMitralyzer (section 7.2) covers the automatic part. The MITK allows to share data among different plug-ins. For the semi-automatic segmentation it is sufficient that the Mitralyzer and SeMitralyzer share an instance of a mitral model (see figure 62). Together with the Mitralyzer, an implementation of an interactor (following the MITK-Interaction-Concept¹) is activated. The interactor represents a state machine that processes user-input depending on its state. For example, if the Mitralyzer view changes to the *Leaflet* tab, the state is changed accordingly to allow interaction with the leaflet points. Depending on the purpose, the manually constructed model can at any point be used as input for the SeMitralyzer and any model produced by the SeMitralyzer can be edited manually.

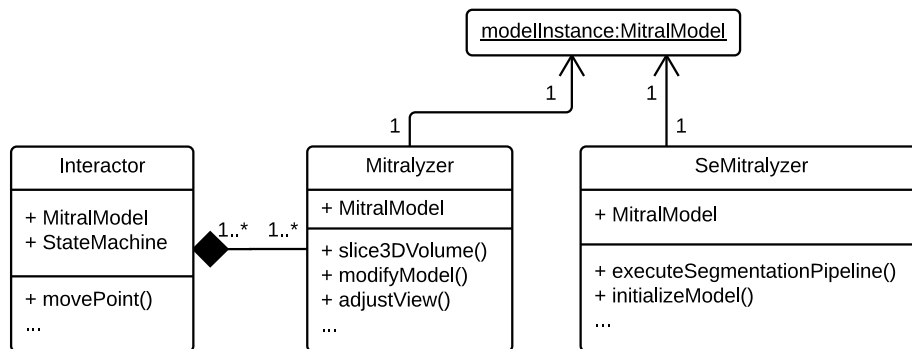


Figure 62: The Mitralyzer and SeMitralyzer are connected via a mitral model instance.

7.2 The SeMitralyzer Plugin

The semi-automatic Mitralyzer, short SeMitralyzer, is the development plug-in that was used to manually execute the segmentation pipeline or single parts of the pipeline. Its root component is a view that gives the user access to all its functionalities. Figure 63 shows the header of the SeMitralyzer GUI. When working on a new segmentation, the user needs to set a model, that represents the input annulus shape, and an image that was used to segment that annulus. The second input and mask input are needed in later steps. A checkbox allows to switch processing of a single time step or the whole data set on or off. This function was only implemented for some of the pipeline components if it was necessary during development.

¹<http://docs.mitk.org/2014.03/DeveloperManualPortal.html>




Figure 63: Header of the SeMitralyzer GUI.

Further, the view offers four tabs. The *Misc* and *Algorithms* tab give access to single components of the pipeline. The *Pipeline* tab is used for execution of the whole automatic segmentation pipeline. The evaluation tool, that will be proposed later in this section, can be started via the *Evaluation* tab. Figure 64 shows the *Algorithms* tab that contains the main components.

Group 1 contains the input fields for the graph cut. The directional weight parameters can be set as well as the source image for the node values. If the option is set to "TTD", the algorithm expects a single input image (usually a TTD image). Another option expects two input images that will then be multiplied and used like the single TTD input. This allows, for example, for modifying the node values with a mask or other weights - anyway, this feature was not used in the final pipeline. The process toggles between 2D and 3D graph cut and the search space sets the graph topology, which can either be a common grid or the space proposed by Schneider. Lastly, Gaussian blur can be applied to the node values.

The second group controls the TTD, where the first parameter should be set to half of the expected leaflet thickness in voxel size and the second one sets the neighbourhood size (refer to section 3.1).

Group 3 is used for the initialization and optimization of the model. The *Active Config* is used to define a location on the hard drive that contains a set of preprocessed data, like graph cut and TTD images, that should be used for the initialization and optimization steps. A checkbox determines if a masked or unmasked graph cut should be used from the available data. The first button executes an initialization approach that is not based on the clustered 3D graph cut, but on the TTD image. This algorithm has not been finished because it did appear not to have enough potential. The next button initializes parts of the model as described in section 6.3.7. Each part of the commissure-coaptation-leaflet initialization can be done separately. The correction of leaflet outliers is optional. Finally, the model optimization process can be executed with the last button.

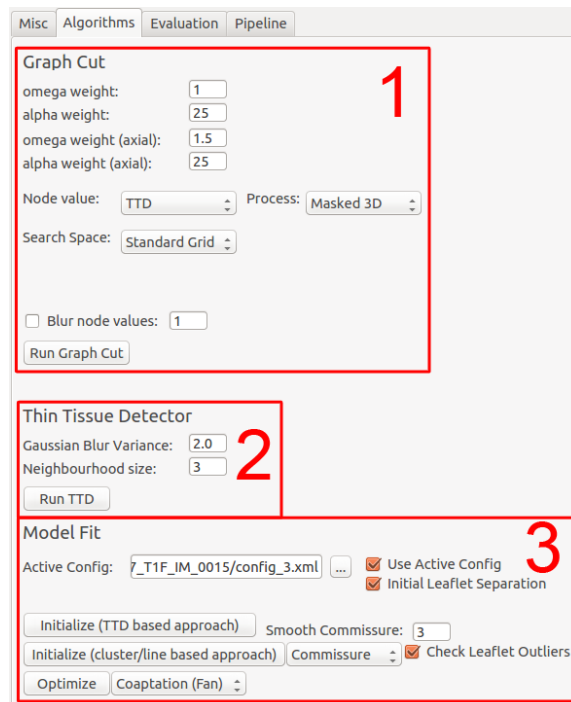


Figure 64: Algorithms Tab.

The *Misc* tab holds buttons to perform simple operations, like mask creation, MIP, image to model alignment and a binary OR to join images. A graph cut can be optimized and separated from the blood pool by either applying a simple threshold or using the clustering described in section 6.3.5.

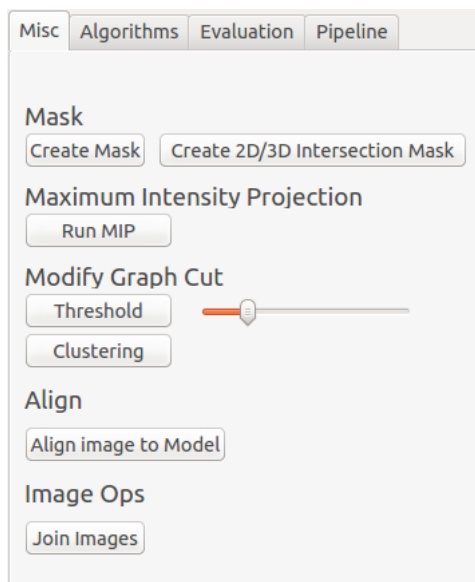


Figure 65: Misc Tab.

All of the components in the *Algorithm* and *Misc* tab require a certain set of input data. The output can be a single or a set of images and is made available to the user via the MITK datamanager-plugin, so that each output can directly be used as input for a next step.

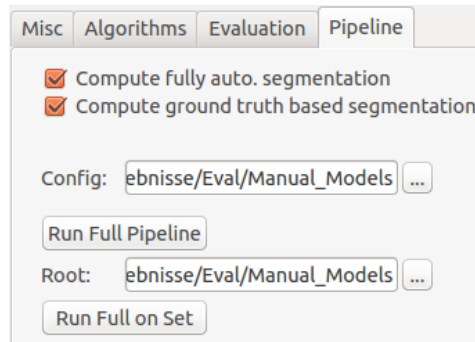


Figure 66: Pipeline Tab.

It is a tedious task to execute each pipeline component one by one. For automation, the *Pipeline* tab can be used. Here, the pipeline from section 6.2 is implemented for execution. The image preprocessing will always be done, but the model generation is optional. The user can choose if a full automatic segmentation should be done and/or a ground truth based segmentation. The ground truth based segmentation takes a ground truth model and recomputes parts of the model as in the *semi-1/2* configuration (refer to section 6.2).

An input annulus model and image can be passed in by the user through the GUI. In this case a config file has to be deployed as well. This file contains all parameters that can be set by the user in the *Algorithm* and *Misc* tab, so it represents a full definition for a pipeline run. Thus, it can be used set to either run the masked or unmasked pipeline (path A or B in figure 34). Alternatively, the user can set a root folder. This folder should contain subfolders, each containing an input annulus model and image and a config file. The Pipeline will then be executed for each subfolder and its specified input.

The result images and models produced by a pipeline run are saved to a folder that has the same location as the config in use and is named after the internal config name. This folder contains data, named by the following structure:

- Images: *ConfigName_Identifier*
- Models: *ConfigName_SegmentationConfiguration_SegmentationStep*

For example, if the internal config name is *Patient01* and the image contains the TTD result for this pipeline, the image would be named *Patient01_TTD*.

The models are separated by their segmentation configuration (*auto* and *semi1-2*) and their last segmentation step. So a file named *Patient01_auto_leaflet* would contain a model that was fully automatically processed throughout the complete pipeline, since the leaflet step is the last one.

The Evaluation Tool can be found in the *Evaluation* tab of the SeMitralyzer plug-in (see figure 67).

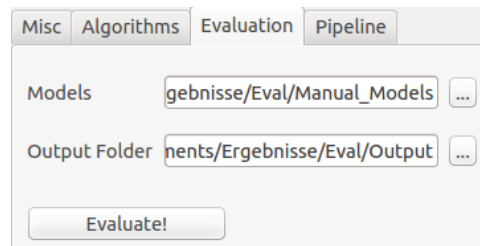


Figure 67: Evaluation Tab.

Inputs are a path that contains all models to be processed and an output folder where the results will be stored. The files in the input folder have to follow a name convention, so that different models can be distinguished:

- *Group-Identifier*

In each group one file must be identified by *groundTruth*. This model will be used as the ground truth for all other models in this group to compare against. The results will be stored in a folder for each group and a file for each *Identifier*. The files contain measurements for all magnitudes that will be proposed in section 8.1.1, for each time step in *.csv* format. Additional measurements can be added, since a *c++* struct is used internally to hold the data, which can be easily extended with more values. Distances that are not just measured point to point, are processed with help of VTK², which is part of the MITK. It offers several methods to obtain the area of a polygon or the distance of a point to a polygon.

Figure 68 shows a class diagram of the SeMitralyzer. The central class is the View. It receives input by the user and executes the functions described above. It has access to an instance of the evaluation tool, a class that wraps the automatic pipeline execution and to each pipeline component. The one to one association to the pipeline components is necessary, because the parameters used in these classes are set in the manner of a state machine. So each instance of a component class should only be accessed by one class. The Pipeline class holds its own pipeline component instances and can execute them sequentially, passing the result of one component to the next one.

²<http://www.vtk.org/>

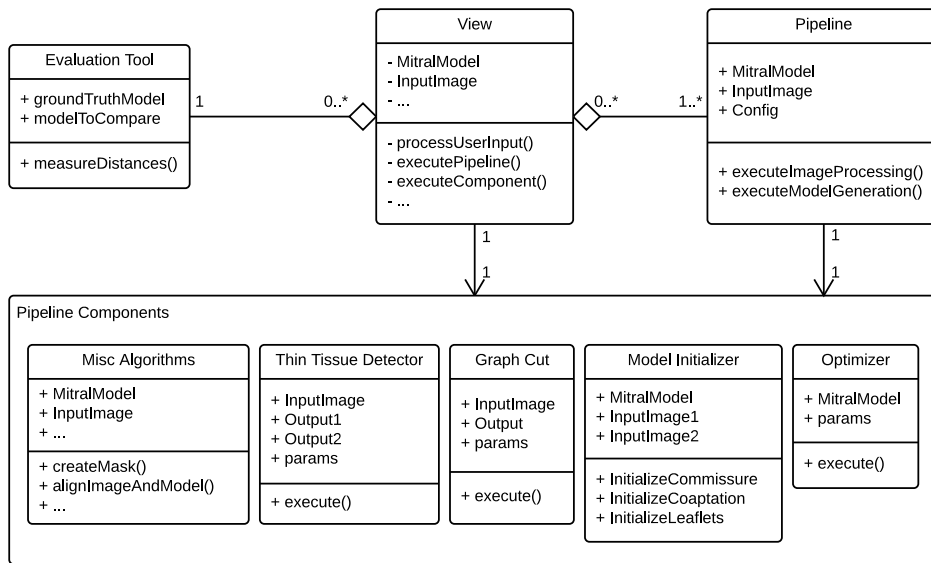


Figure 68: SeMitralyzer class diagram.

7.3 Mitral Util

The Mitral Util was first implemented by Graser [6] to allow a visual and quantitative examination of annulus rings over all available time steps. It has now been augmented with some of the magnitudes that will later be used for evaluation and with some prototypical measurements that show how the mitral model with leaflets might give opportunities for new medical research topics. The new displayed magnitudes are:

- Orifice Area 2D/3D
- Anterior/Posterior Leaflet Area
- Orifice Circumference
- Coaptationarea
- Annulus to Orifice Angle

For each magnitude the user is given information about the average, maximum, minimum and dynamic range value of all time steps and the value of the currently selected time step (see figure 69). The Annulus to Orifice Angle (AOA) is an example for a measure that could not be investigated without a model of the leaflets. It describes the angle between the annulus plane normal and the orifice plane normal. The orifice plane is a plane that fits into the coaptation points of a mitral model.

	Current	Avg	Min	Max	Dynamic	
O. Area 2D	412.961	231.281	42.3175	736.291	693.974	mm ²
O. Area 3D	622.796	315.412	62.2787	993.738	931.459	mm ²
O. Area 2D/3D	.663076	.713901	.656742	0.81692	.160178	

Figure 69: Display section of the Mitral Util showing some of the available measurements.

The Mitral Util might be used for a survey about the correlation between a certain AOA and certain pathologies of the MV. The other magnitudes were actually used in this work and are described in more detail in section 8.1.1. Figure 70 shows an example plot of the leaflet area over all time steps of a mitral model. The data was taken from a manual segmentation. With such plots, a segmentation can be tentatively checked for reasonable results.




Figure 70: Example plot of the leaflet area over time.

8 Experiments and Results

The quality of the segmentation proposed in this work has to be measured. This section will describe the evaluation process and present the results. The experiments part introduces several quantifications and explain their relevance. The result section will compare results of the automatic segmentation with a manually segmented ground truth and show how much the semi-automatic segmentation can improve a manual work-flow with respect to segmentation time.

8.1 Experiments

Other modelling approaches presented in the related work section have used a ground truth model manually segmented by an expert. This ground truth was then used to compare their (semi-)automatic models against. In this work the same is done to determine the quality of the leaflet segmentation. An expert has segmented five MVs over a full cardiac cycle, with

an overall of 90 time steps (43 at systole and 47 at diastole). The Live-3D mode of the Philips ultrasound system iE33 xMatrix and a Philips X7-2t matrix array transducer (Philips Healthcare, Andover, MA, USA) were used to obtain the image data. The annulus rings of these manual segmentations were used as the input for the automatic segmentation, since the semi-automatic annulus segmentation proposed by Graser is not part of the evaluation, even though it is part of the mitralyzer plug-in. The semi-automatic segmentation should be quantified as well. For this purpose a set of MVs was segmented with or without support of the automatic initialization. This way the time saved of a semi-automatic segmentation can be measured compared to a fully manual segmentation.

The next section introduces the magnitudes that are measured during the experiments and the subsequent section briefly describes the evaluation process.

8.1.1 Magnitudes

First of all the magnitudes used for the evaluation will be described in detail. The following list summarizes the measurements done to assess the automatic segmentation:

- Model to Model Distance - point to surface $D_s(A, B)$ (equation 9)
- Model to Model Distance - Hausdorff $H(A, B)$ (equation 10)
- Commissure to Commissure Distance $Comm(A, B)$ (equation 11)
- Coaptation to Coaptation Distance $Coap(A, B)$ (equation 12)
- Leaflet Area $L(A)$
- Coaptation Area $C(A)$ (equation 13)
- Orifice Area $OA(A)$
- Orifice Circumference $OC(A)$

Most related work has used a model to model distance to measure their segmentation results. In this work this is done by using a point to surface distance measurement. When comparing model A and B, the shortest distance of every point in A to the surface of B is summed up. $D(A, B)$ is the mean distance of model A to B.

$$D(A, B) = \frac{\sum_{a \in A} d(a, B)}{|A|}$$

where $d(a, B)$ is the minimum distance of point a to the surface of B . It is not guaranteed that $D(A, B) = D(B, A)$, so a symmetric distance is defined as

$$D_s(A, B) = \frac{D(A, B) + D(B, A)}{2} \quad (9)$$

The requirement of a symmetric distance gets clear when looking at figure 71, illustrating the problem in 2D. It shows a curve A and B and some exemplary distances from A to B and B to A. It is obvious that the high peak in B is skipped in $D(A, B)$, because a closer point on B can always be found. This does not happen in $D(B, A)$.




Figure 71: Comparison of $D(A, B)$ and $D(B, A)$.

Another distance measure that considers how well two models cover is the Hausdorff distance [36]. The original definition is that it "computes the maximum distance of a set to the nearest point in the other set" [37]. Instead of comparing the sets of the model points, the distances can again be computed by the point to surface function $d(a, B)$. Symmetry has to be taken into account as well and the Hausdorff distance $H(A, B)$ in this work is defined as

$$H(A, B) = \max \left\{ \max_{b \in B} [d(b, A)], \max_{a \in A} [d(a, B)] \right\} \quad (10)$$

The quantifications so far are very general, the following ones test the segmentation for more detailed clinical relevance. The commissure points determine the width of the open valve as well as the length of the coaptation line, so it is helpful to know how good the automatic segmentation matches the ground truth for these points. The assessment is very simple as merely the Euclidean distance from a commissure point in one model to the affiliated point in the other model is taken.

$$Comm_{al}(A, B) = |a_{al} - b_{al}| \text{ and } Comm_{pm}(A, B) = |a_{pm} - b_{pm}| \quad (11)$$

where x_{al} is an antero-lateral and x_{pm} a postero-medial commissure point.

The coaptation to coaptation distances $Coap_a(A, B)$ and $Coap_p(A, B)$ state how well the anterior or posterior leaflet tips match in model A and B. This distance is tailored to the coaptation optimization process of section 6.3.8.

Remember how the coaptation points were able to move on their respective fan slices. Since the fan slicing range depends on the location of the commissure points and the commissure points can vary among the models to be compared, a space must be found that includes the coaptation lines of both models. The illustration in figure 72 shows how this space is constructed. The view direction in the left and center image is the annulus normal. The right image shows a side view. The outer slices are defined so that they touch both coaptation lines and at least one of them at the start/end point and run through the saddle horn. The dashed part of each coaptation line is cut off. The remaining parts are subdivided into an equal number of points.




Figure 72: Illustration of the fan slice range that covers both models' coaptation lines (gray area). A/B: anterior/posterior coaptation of model A/B. The dashed lines in the left image are parts of the coaptation lines that are not considered for the measurement. a and b: Subdivision points of the respective curve. Right: Example slice, constructed from the dashed line in the center image.

Now the distance is taken from each point in A to its affiliated point in B. The mean of all distances is the coaptation distance.

$$Coap_x(A, B) = \frac{\sum_{i=1}^n |a_i - b_i|}{n}, a \in A_x, b \in B_x \quad (12)$$

where n is the number of subdivisions and x can either be a or p to represent the measurement for the anterior or posterior coaptation line. A_x and B_x are the subsets of points that express the line. This method assumes that, because of a high n , the point pairs that are compared are almost within the same fan slice. The re-sampling is required because the segments of the coaptation lines that are taken into account can vary and assures an equal number of points in both coaptation lines. In this case the distance is the offset that a point would have had to move within its slice during the optimization process, to reach to the location proposed by the ground truth (refer to the right image in figure 72).

An absolute magnitude is the leaflet area. It may be of interest to a surgeon

to know how much the leaflet area varies over time. When anatomical assumptions are made it can also be used to qualify a manual segmentation, for example if the expected maximum area variety is at 35% in radial direction as in [38]. The leaflet area of a valve A will be denoted as $L_a(A)$ for the anterior and $L_p(A)$ for the posterior leaflet.

A value that, to our knowledge, has not yet been measured in any publication of automated MV modelling is the coaptation area, which is the surface of the leaflets that actually coapt during valve closure. Whether the MV closes watertight or not depends on the coaptation area, so this magnitude might be of the highest clinical relevance of all magnitudes proposed in this work. To compute the coaptation area, the point to model distance $d(a, B)$ is used, but in this case the points are taken from one leaflet and the surface is taken from the other leaflet of the same valve model. This is done in both directions. All points that are closer to the other leaflet than a threshold are taken into account for the area calculation. Let A_a be the point set representing the anterior leaflet and A_p the posterior leaflet. Then,

$$A_{ac} \subseteq A_a \wedge \forall x \in A_{ac} : d(x, A_p) \leq t \wedge \forall y \in (A_a - A_{ac}) : d(y, A_p) > t$$

$$A_{pc} \subseteq A_p \wedge \forall x \in A_{pc} : d(x, A_a) \leq t \wedge \forall y \in (A_p - A_{pc}) : d(y, A_a) > t$$

$$C(A) = \frac{Area(A_{ac}) + Area(A_{pc})}{2} \quad (13)$$

where t is the expected leaflet thickness or maximum distance and $Area(x)$ computes the surface area of all connected vertices in x . Figure 73 demonstrates the concept. The solid black lines are the leaflet surfaces. The gray area surrounding them is defined by the expected leaflet thickness, which was set to 2 mm in this work. The gray areas touch within the dashed rectangle. Here, the leaflet surfaces have a distance of at most the expected leaflet thickness.




Figure 73: Schematic of the coaptation area measurement.

The coaptation area could be further analysed to find regurgitant orifices during valve closure, as figure 74 shows. The coaptation surface is shown

as a yellow surface, the rest of the valve is transparent. It can be observed that the coaptation surface has a gap in the marked region. This gap could be interpreted as an orifice. The available data could be used to detect such gaps automatically, but this was not implemented in this work.

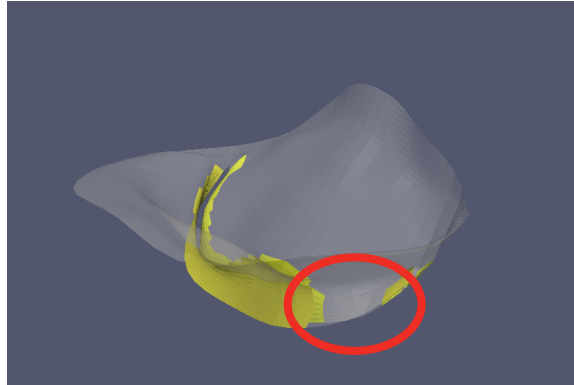


Figure 74: Visualization of the coaptation area. The marked region shows a gap in the coaptation area.

The last two magnitudes are the orifice area $OA(A)$ and the orifice circumference $OC(A)$. All coaptation points of both leaflets are connected to a 3D surface. $OA(A)$ is the area of this surface. $OC(A)$ is the added length of both coaptation lines. These values might be of interest when comparing pre- and post-operative valves, after a resection has been done.

8.1.2 Process

This section will describe the process that was carried out to create models for the automatic and semi-automatic evaluation, as well as for to measure the time saved compared to a manual segmentation.

As mentioned earlier, an expert has segmented 5 data sets over a full cardiac cycle. There were more data sets available, but the image quality was not sufficient for the automatic algorithm (refer to section 6.3.3, about US artifacts). Among the problems of the unused images were heavy artifacts, a disadvantageous field of view so that parts of the valve were not visible in some time steps, or a very low and varying contrast so that it was difficult to distinguish tissue from the blood pool.

Those five data sets were then (semi-)automatically segmented using both methods, with and without a 3D graph cut based on leaflet separation (path A and B in figure 6.2). The distance measurements of each result model to its ground truth were then computed. For the absolute magnitudes the ground truth value was subtracted from the result value to find out the direction and extent of the deviation. The semi-automatic segmentation was

done from two different starting configurations. Together with the full automatic process there are three configurations:

1. No user interaction. (*auto*)
2. The expert has set both commissure points. (*semi-1*)
3. The expert has set commissure and coaptation points. (*semi-2*)

The goal was to find out how much user interaction is required to lead the semi-automatic approach to a reliable result. The expert has labelled each time step as diastole or systole, so the results can be averaged for each valve state which might unveil state related performance differences.

To obtain the time saved by the semi-automatic approach, it was measured how long it takes to segment a time step, averaged from 106 manually segmented time steps. Measures were taken at three stages:

1. Time to set the commissure points.
2. Time to set the coaptation points.
3. Time to set the leaflet points.

A portion was done completely manually to determine an average time value without support. For the remaining time steps, the expert got proposed a solution for each step as described in section 6.2. So the points had to be checked and relocated if necessary. The results will show in which stages the user had the greatest benefit of the semi-automatic support.

8.2 Results

This section will state the results obtained during the evaluation process and is divided into three parts. The first one deals with the (semi-)automatic segmentation results and the second one focusses on the manual segmentation time evaluation. A runtime evaluation for the automatic pipeline can be found in the third part.

An example for an automatic and a manual segmentation of the same data set is displayed in figure 75. Figure 76 shows the automatic segmentation results over a whole cardiac cycle.




Figure 75: Segmentation result example. Left: Automatic segmentation. Right: Manual segmentation.




Figure 76: Automatic segmentation over a whole cardiac cycle.

8.2.1 Semi-Automatic Segmentation

As a first step the overall performance of the masked and unmasked pipeline will be confronted. Table 5 shows the averaged results of the full automatic segmentation and the two semi-automatic configurations for both, the masked and unmasked pipeline. The values printed in boldface are the better values of each magnitude. It can be observed that the masked pipeline is slightly superior in 7 and inferior in 3 of 11 cases. Anyway, the differences are so small that they can almost be neglected. Considering the additional computation time that the masked approach requires (see section 8.2.3), the unmasked approach is the pipeline of choice. Consequently, the remaining evaluation will only focus on the unmasked pipeline.

	D_s	H	$Comm_{al}$	$Comm_{pm}$	$Coap_a$	$Coap_p$	ΔC
masked	0.68	4.46	5.37	4.50	4.26	5.22	19.44
unmasked	0.73	4.77	5.37	4.50	4.34	4.30	20.61
	$D_{s,a}$	$D_{s,p}$	ΔL_a	ΔL_p	ΔOA	ΔOC	/
masked	0.79	0.93	76.61	85.54	33.93	8.94	/
unmasked	0.86	0.97	74.49	84.91	33.44	9.00	/

Table 5: Direct comparison of masked and unmasked overall performance. All values given in *mm*.

The averaged results of table 5 were composed of the full automatic (*auto*) and the two semi-automatic configurations - one with user input for the commissure (*semi-1*) and one with the user input for the whole coaptation (*semi-2*). The values for the commissure distance were taken from *auto* only, because the commissure is placed manually in both other cases. The same applies to the coaptation distance, which is only valid for *auto* and *semi-1*. The plot in figure 77 shows the symmetric point to model distance for the whole model and for each leaflet. It should be mentioned that at first the time steps of each model were averaged to assign one value per magnitude to one model. Then the values of all models were averaged again to obtain an overall value for each magnitude. Single time steps of models with more time steps have less weight for the final result in this model-based evaluation.




Figure 77: The plot shows the symmetric point to model distance and Hausdorff distance for the whole model and for each leaflet.

The leaflets can be further subdivided into three segments each. The following tables show the symmetric distance for the whole valve and the six segments. The Hausdorff distance is also displayed for reference. The values shown are the average over all, systolic and diastolic time steps.

Note that the mean is given by averaging all time steps, not the individual results for each model. This time-step-based evaluation causes slightly different results compared to model-based evaluation in figure 77 and table 5. This is done because every time step should be treated individually, since there is no 4D information used except for the outlier correction. The standard deviation is shown as well as the 95th percentile of each set, meaning that 95% of all time step measures resided below that value. Better values of diastole or systole are highlighted boldfaced.

	D_s	$D_{s,A1}$	$D_{s,A2}$	$D_{s,A3}$	$D_{s,P1}$	$D_{s,P2}$	$D_{s,P3}$	H
A \emptyset	0.82	1.27	1.19	1.08	1.56	1.35	1.23	5.31
A σ	0.25	0.43	0.43	0.43	0.48	0.56	0.46	1.94
A 95%	1.31	2.24	2.42	2.04	2.66	2.53	2.35	10.2
D \emptyset	0.91	1.26	1	1.05	1.53	1.09	1.15	5.93
D σ	0.24	0.42	0.3	0.47	0.48	0.48	0.49	1.92
D 95%	1.34	2.08	1.87	2.03	2.59	2.19	2.25	10.07
S \emptyset	0.73	1.28	1.4	1.12	1.6	1.63	1.31	4.63
S σ	0.21	0.45	0.49	0.4	0.49	0.51	0.43	1.78
S 95%	1.24	2.28	2.6	1.98	2.68	2.73	2.33	10.07

Table 6: Symmetric and Hausdorff distance overview for *auto* segmentation. All time steps (A) and systolic (S) and diastolic (D) subsets for anterior (A1, A2, A3) and posterior (P1, P2, P3) leaflet segments. All values given in *mm*.

The segmentations for each separate leaflet segment are better during diastole in table 6. Systolic values are better when considering the whole valve as in D_s and H .

	D_s	$D_{s,A1}$	$D_{s,A2}$	$D_{s,A3}$	$D_{s,P1}$	$D_{s,P2}$	$D_{s,P3}$	H
A \emptyset	0.8	0.88	1.11	0.88	0.99	1.38	0.83	5.69
A σ	0.23	0.39	0.42	0.33	0.42	0.61	0.31	2.47
A 95%	1.2	1.52	2.51	1.58	1.86	2.84	1.59	12.16
D \emptyset	0.86	0.93	0.9	0.85	1	1.1	0.75	6.19
D σ	0.22	0.41	0.23	0.36	0.45	0.54	0.37	2.58
D 95%	1.2	1.62	1.45	1.7	1.81	2.48	1.82	13.17
S \emptyset	0.74	0.82	1.34	0.91	0.98	1.68	0.91	5.15
S σ	0.2	0.36	0.56	0.3	0.4	0.62	0.24	2.26
S 95%	1.19	1.39	2.62	1.45	1.88	2.98	1.42	11.35

Table 7: Symmetric and Hausdorff distance overview for *semi-1* segmentation. All time steps (A) and systolic (S) and diastolic (D) subsets for anterior (A1, A2, A3) and posterior (P1, P2, P3) leaflet segments. All values given in *mm*.

For the results in table 7 the commissure points were given. The values have improved and the quality difference between diastole and systole has lessened.

	D_s	$D_{s,A1}$	$D_{s,A2}$	$D_{s,A3}$	$D_{s,P1}$	$D_{s,P2}$	$D_{s,P3}$	H
A \emptyset	0.45	0.53	0.43	0.51	0.53	0.39	0.47	2.07
A σ	0.23	0.35	0.28	0.34	0.31	0.23	0.27	0.8
A 95%	0.89	1.31	0.94	1.28	1.19	0.87	1.02	4.01
D \emptyset	0.5	0.61	0.46	0.59	0.58	0.34	0.51	2.31
D σ	0.28	0.45	0.31	0.41	0.37	0.22	0.31	0.88
D 95%	0.95	1.38	0.94	1.37	1.41	0.68	1.03	4.19
S \emptyset	0.4	0.45	0.41	0.43	0.48	0.42	0.42	1.81
S σ	0.18	0.24	0.24	0.24	0.23	0.2	0.22	0.62
S 95%	0.71	0.85	0.86	1.02	0.91	0.92	0.95	3.24

Table 8: Symmetric and Hausdorff distance overview for *semi-2* segmentation. All time steps (A) and systolic (S) and diastolic (D) subsets for anterior (A1, A2, A3) and posterior (P1, P2, P3) leaflet segments. All values given in *mm*.

With the whole coaptation given, the results in table 8 show that systolic time steps have now been processed more successfully.




Figure 78: The plot shows the symmetric point to model distance for each leaflet segment in each configuration.

The plot in figure 78 reveals that the diversity among the segments' distance values decreases with increasing user input. The best performing segment is different for almost all configurations. Anyway, a trend is visible that A3 and P3 perform good with most stability.

The following tables contain the results of the other magnitudes for each configuration. The *average percent difference* (apd) is a relative measure, showing how much the mean error diverges from the mean of the ground truth absolute value. This is helpful since the absolute values of some quantification measurements change naturally over the cardiac cycle. For example, if a mean ground truth area is 10 mm^2 and the mean segmented area is 6 mm^2 , then $\text{apd} = 40\%$.

Better values of diastole and systole are highlighted boldfaced and anterior/posterior comparisons are highlighted in italics. This is done for the mean values of $Comm$, $Coap$, ΔL and ΔOC and for apd of the differences of the remaining magnitudes. It is necessary to look at the apd in these cases, since the absolute values may change severely from diastole to systole. For example the orifice area is near zero in closed valve state and reaches almost up to 1000 mm^2 in open valve state. A similar case applies to the coaptation area.

	$Comm_{al}$	$Comm_{pm}$	$Coap_a$	$Coap_p$	ΔC	ΔL_a	ΔL_p	ΔOA	ΔOC
A <i>apd</i>	/	/	/	/	58.87	15.35	17.25	20.73	15.03
A \emptyset	5.1	4.45	4.98	5.23	23.73	<i>116.88</i>	123.83	50.36	16.03
A σ	2.92	2.75	1.74	1.9	23.73	111.12	117.8	41.68	13.91
A 95%	11.89	10.08	9.35	9.79	74.39	396.11	302.57	135.22	38.6
D <i>apd</i>	/	/	/	/	52.54	10.59	14.03	14.33	10.88
D \emptyset	4.15	3.62	4.42	4.17	21.18	81.21	106.27	52.64	12.17
D σ	2.37	1.96	1.47	1.55	22.64	80.85	102.34	47.6	12.29
D 95%	10.18	8.48	8.48	7.69	94.9	188.19	311.07	150.19	33.85
S <i>apd</i>	/	/	/	/	47.54	20.61	<i>21.21</i>	44.8	20.04
S \emptyset	6.13	5.36	5.59	6.4	26.52	155.87	<i>143.01</i>	47.86	20.24
S σ	3.39	3.27	1.99	1.83	26.04	118.99	133.41	36.4	10.74
S 95%	13.42	11.15	10.11	10.72	56.28	415.56	297.48	123.15	40.14

Table 9: Result overview for *auto* segmentation. All time steps (A) and systolic (S) and diastolic (D) subsets. Distances given in mm , areas given in or mm^2 , apd given in %.

It can be observed that the results for the open valve state are superior in most cases. The anterior leaflet area was predicted better than the posterior on average. The plots in figure 79 show detailed results for a selected data set.




Figure 79: The plots show the coaptation and orifice area of the *auto* segmentation and ground truth of a selected data set over time. The vertical line in the right plot shows where the user has defined the transition between systole and diastole.

	$Coap_a$	$Coap_p$	ΔC	ΔL_a	ΔL_p	ΔOA	ΔOC
A <i>apd</i>	/	/	55.89	14.17	15.74	19.09	10.72
A \emptyset	4.05	4.65	22.53	107.92	113.01	46.37	11.44
A σ	1.88	1.81	22.45	89.5	97.96	46.01	10.95
A 95%	8.6	9.1	67.57	339.19	242.2	132.68	29.21
D <i>apd</i>	/	/	37.31	8.03	13.06	37.31	6.49
D \emptyset	3.2	3.58	15.04	61.53	98.93	15.04	7.26
D σ	1.36	1.24	14.82	56.73	87.47	14.82	7.33
D 95%	7.63	6.86	54.7	172.01	277.1	54.7	25.69
S <i>apd</i>	/	/	55.06	20.98	19.04	55.06	15.85
S \emptyset	4.99	5.81	30.72	158.64	128.39	30.72	16.01
S σ	2	1.8	30.8	103.66	109.51	30.8	12.98
S 95%	9.15	9.34	89.24	349.43	236.96	89.24	31.44

Table 10: Result overview for *semi-1* segmentation. All time steps (A) and systolic (S) and diastolic (D) subsets. Distances given in *mm*, areas given in or *mm*², *apd* given in %.

For the *semi-1* segmentation, the better results again reside within the open valve time steps (table 10).

	ΔC	ΔL_a	ΔL_p
A <i>apd</i>	23.47	2.04	2.08
A \emptyset	9.46	15.5	14.92
A σ	10.2	15.43	14.73
A 95%	44.48	47.67	52.66
D <i>apd</i>	10.54	2.15	1.68
D \emptyset	4.25	16.51	12.73
D σ	4.47	15.68	12.8
D 95%	12.09	48.88	50.15
S <i>apd</i>	27.16	1.9	2.57
S \emptyset	15.15	14.4	17.32
S σ	15.89	14.59	16.19
S 95%	70.83	39.75	43.03

Table 11: Result overview for *semi-2* segmentation. All time steps (A) and systolic (S) and diastolic (D) subsets. All values given in mm^2 , *apd* given in %.

Table 11 shows that, given a complete coaptation line, the leaflets can be found with more similar precision, whereas the coaptation is better predicted during diastole.

Tables 12 and 13 show how much the results of the *auto* segmentation were improved by *semi-1* and *semi-2* respectively. The numbers represent the percentage of improvement compared to the *auto* values.

$$val = (val_{auto} - val_{semi}) / val_{auto} * 100$$

	$Coap_a$	$Coap_p$	ΔC	ΔL_a	ΔL_p	ΔOA	ΔOC
A <i>apd</i>	/	/	5.73	7.69	8.75	7.91	28.68
A \emptyset	18.67	11.09	5.06	7.67	8.74	7.92	28.63
A σ	-8.05	4.74	5.39	19.46	16.84	-10.39	21.28
A 95%	8.02	7.05	9.17	14.37	19.95	1.88	24.33
D <i>apd</i>	/	/	28.99	24.17	6.91	15	40.35
D \emptyset	27.6	14.15	28.99	24.23	6.91	14.99	40.35
D σ	7.48	20	34.54	29.83	14.53	6.01	40.36
D 95%	10.02	10.79	42.36	8.6	10.92	14.04	24.11
S <i>apd</i>	/	/	-15.82	-1.8	10.23	-0.6	20.91
S \emptyset	10.73	9.22	-15.84	-1.78	10.22	-0.61	20.9
S σ	-0.5	1.64	-18.28	12.88	17.91	-20.66	-20.86
S 95%	9.5	12.87	-58.56	15.91	20.34	-7.26	21.67

Table 12: Improvement achieved by *semi-1* segmentation. All values given in %.

Considering all values equally, the *semi-1* segmentation has achieved an overall improvement of 11.12%.

	ΔC	ΔL_a	ΔL_p
A <i>apd</i>	60.13	86.71	87.94
A \emptyset	60.13	86.74	87.95
A σ	57.02	86.11	87.5
A 95%	40.21	87.97	82.6
D <i>apd</i>	79.94	79.7	88.03
D \emptyset	79.93	79.67	88.02
D σ	80.26	80.61	87.49
D 95%	87.26	74.03	83.88
S <i>apd</i>	42.87	90.78	87.88
S \emptyset	42.87	90.76	87.89
S σ	38.98	87.74	87.86
S 95%	-25.85	90.43	85.54

Table 13: Improvement achieved by *semi-2* segmentation. All values given in %.

Considering all values equally, the *semi-2* segmentation has achieved an overall improvement of 75.21%.

The next section will be about the user effort that is necessary to segment a valve (semi-)manually, to show the trade-off between segmentation time and quality.

8.2.2 Manual Segmentation Time Measure

To obtain the required user effort to segment a MV manually or semi-automatically, the time that it took an expert to segment one time step was measured for a total of 103 time steps. 35 time steps were done completely manually. The remaining time steps were done with aid of the automatic algorithm, as shown in figure 36. Checkpoints were taken after commissure-, coaptation- and leaflet-segmentation.

	Commissure	Coaptation	Leaflets	Total
w/o proposed points	68.44	197.81	221.25	487.5
with proposed points	57.76	109.25	10.59	177.6
improvement	15.6%	44.77%	95.21%	63.57%

Table 14: Total segmentation time in seconds and time improvement.

An expert can segment a valve with given annulus in 487.5 seconds using the manual segmentation method proposed in section 6.1. With support of the automatically proposed points, the time reduces to 177.6 seconds (see table 14). The leaflet segmentation benefits the most of the proposed segmentation. It has to be considered that the manual commissure and coaptation segmentation time can vary significantly, depending on the image quality and MV anatomy. Thus, assuming that the commissure and

coaptation improvement was not improved by the automatic support, but by a MV easier to delineate, a total segmentation time with leaflet proposal would still be 276.84 seconds, with an improvement of 43.21%.

8.2.3 Runtime Analysis

The runtime of the segmentation pipeline and its single components was measured while processing the five data sets that were used for evaluation on a machine with four 2.4 GHz cores and 8GB RAM. Table 15 shows the total runtime for volumes with a size of $224 \times 208 \times 224$ voxels per time step. The pipeline was partially run with multi-threading, processing four time steps at once. If less time steps are available, the processing time per time step increases, since the parallelism can not be exploited. The time required for a single time step without multitasking is shown as well.

	avg.	max	min
(p) masked	67.93	88.87	52.23
(p) unmasked	31.3	38.16	23.40
(s) masked	100.89	132.41	85.78
(s) unmasked	35.96	39.06	33.09

Table 15: Total runtime in seconds per time step for masked and unmasked pipeline with partial multi-threading (p) and sequential execution (s). All values given in seconds.

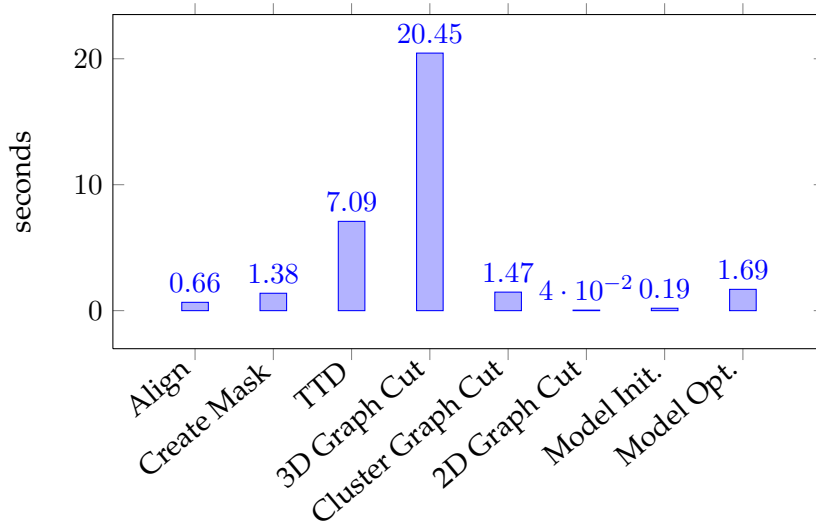


Figure 80: Plot showing the runtime of each pipeline component of the unmasked pipeline in seconds per time steps for sequential execution.

The runtime of each component of the pipeline was analysed, too. The plot in figure 80 shows the times of the unmasked pipeline. For the masked

pipeline two additional 3D graph cuts would be necessary (refer to the masked pipeline in section 6.2), together with some preparation. This took on average 65.26 additional seconds.

The 3D graph cut occupies the largest portion of the processing time with 56.87%, the TDD takes 19.72%, the other components can almost be neglected on this scale.

9 Discussion

This section will discuss and interpret the results from section 8 and the observations done during this work. Section 9.1 deals with the measurement results, the manual modelling technique and ground truth generation will be examined in 9.2. Finally, the semi-automatic approach will be discussed in 9.3.

Generally, the evaluation done in this work is comparable with the methods summarized in section 4.2.5. The number of available time steps in these works ranges from four to 1516. This work uses 90 time steps and a manually segmented ground truth model is available for each of them.

9.1 Measurements

The overall results have shown that the segmentation proposed in this work is comparable with existing work in this field. It achieved a model to ground truth distance of 0.82 mm and Hausdorff distance of 5.84 on average for the *auto* segmentation considering all single time steps, while the mentioned related work in section 4 reached a mean error of 0.76-1.5 mm and 5.57 mm Hausdorff distance. Anyway, on the one hand this does not mean that the approach in this work outperforms the others. A huge difference is, that here, the annulus shape was given and used for both, segmentation and ground truth, whereas the related methods did not have this advantage. Consequently, a direct comparison is not possible. On the other hand, the Hausdorff distances are similar, so that it can be assumed that this approach is not worse than the others.

The US images used in this work cover approximately $1.4mm^3$ per voxel, so the model to model distance suggests that the algorithm precision is within the image precision. Nevertheless, a visual examination as in figure 81 can easily uncover that a set of more specific measurements is necessary to justify a good segmentation.

The new magnitudes proposed in 8.1.1 were only possible because of the semantic information given by the labelled commissure and coaptation points.

D_s can be measured for every point on a surface and can be intuitively visualized as in figure 81. The image unveils a great weakness of D_s . It is easily visible that,

1. the ground truth model is in open valve state, the automatically segmented model in closed state.
2. the coaptation line is overall shifted to the top of the image in the automatic model.
3. the lower leaflet of the automatic model reaches into the upper leaflet of the manual model, leading to good distance values.




Figure 81: Left: Ground truth model. Right: Automatic segmentation with distances.

Regardless of these three facts, the worst distances are at about 3 mm, as the color code shows. The conclusion of this observation is that it is very important to compare the leaflets separately and that D_s is not precise enough to yield a reliable evidence for the quality of a segmentation. Above all, it does not qualify for clinically relevant use.

As tables 6,7, 8 and the plot in figure 77 show, D_s for the whole valve is lower than the average of D_s of all separate segments or leaflets. The reason for this is that, for example, the anterior leaflet of a model can overlap with the posterior leaflet of its ground truth, which may lead to a lower distance. Figure 82 shows a schematic slice through both leaflets of a model and its ground truth. It is obvious that the segmentation in this example is not accurate, because the orifice is shifted significantly. But still, in the global case, D_s would yield a good result, since one leaflet of a model can be compared to the wrong leaflet of the other model. This cannot happen in the leaflet-specific case, leading to greater distance values.

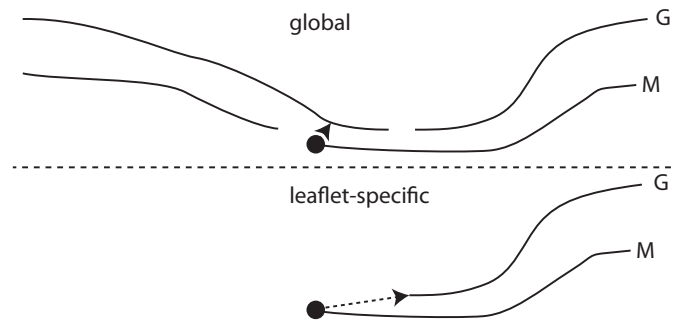


Figure 82: Difference between global and leaflet-specific D_s . Model (M) and its ground truth (G).

The difference between the whole valve and single segment measurement is significant, as the best segment of the *auto* segmentation has a mean error of 2.04 mm and worst segment an error of 2.66 mm. A similar observation can be made for the *semi-1* configuration, though the differences are less severe, since the overlapping issue is confined by the user defined commissure points. The effect is completely lost in the *semi-2* segmentation, because the coaptation is given by the user and an overlapping as in figure 82 is impossible.

An overview over all segment measures has shown that the A3 and P3 segment were estimated best on average. Why this region of the valve was done better is unclear and might be answered by a physician, a possible cause might be the point of view of the US probe. In addition, diastolic time steps were segmented with less error with the *auto* configuration, whereas the systolic time steps did better in *semi-2* configuration. The reason for this is probably, that it is easier for the algorithm to find the leaflet tips in diastole instead of finding the coaptation line in systole.

A similar observation was made in [12], where parts of the valve were labelled as anterior and posterior leaflet and the labelled mean error was higher than the unlabelled error. Also, their unlabelled error was smaller during systole, which is the same as in table 6.

Another conclusion can be based on the Hausdorff distance. It decreases from 5.15 mm, if the coaptation is not given, to 1.81 mm if the coaptation is given, so that the coaptation points are most likely the ones that contribute most to a high Hausdorff distance and are the most difficult to detect. Further, it is obvious that *semi-2* does best, since it starts with a coaptation line defined by the user. In this case $D_s \approx D_{s,a} \approx D_{s,p}$ because the leaflets cannot overlap. The single leaflets can only vary in their surface curvature. The high values for *auto* and *semi-1* are a hint to an unsatisfactory automatic detection of the coaptation points, while *semi-1* is slightly better than no user input.

A look at the general model difference in figure 77 has shown that the an-

terior leaflet segmentation performed slightly better than the posterior one.

While a model to model distance is present in most related work, an investigation on how good specific landmarks were estimated is rarely found. In this work several magnitudes were proposed that use landmark information and focus on clinical relevance. The ability to (semi-)automatically model a MV and then measure, for example, the coaptation area could be used in medical interrogations or for the comparison of pre- and post-operative valves, without requiring too much labour time. As the evaluation revealed, such specific measures are also crucial to determine the quality of a (semi-)automatic segmentation. Tables 9, 10 and 11 show the results for those new magnitudes. The mean values do not yet tell much about the performance, but with the aid of the apd-value it is easy to see that it is way harder to achieve good values for those measures than for the model to model distance.

The commissure to commissure distances for all time steps are on average 4.78 mm. This is ca. 13% of the average valve diameter found in [6], which is undoubtedly a lot, considering that the space between both commissure points influences the maximum possible blood flow, and the maximum orifice area, respectively. 5% of the time steps even were displaced by more than 10.99 mm.

An average offset achieved by *auto* of 5.11 mm for the coaptation lines is also about 15% of the valve size. The commissure points as well as the coaptation points were predicted with less error during diastole. This may again be caused by the fact that the open valve edges are easier to detect. Particularly the voxel classification that is used during the model initialization works best if there is enough actual blood pool included in the considered volume. In the *semi-1* configuration the coaptation offset reduces slightly to 4.35 mm, improving it by 14.87% on average. Overall, the anterior coaptation improves most when the commissure points are user-defined.

The coaptation area is one of the values with most clinical relevance, because it can be used to determine how well a valve closes. Unfortunately, the *auto* segmentation has not done very good to predict it correctly. An apd of 52.54% was achieved during diastole and 47.54% during systole. *Semi-1* has done a bit better, with an improvement of 28.99% during diastole, but with a systolic error increased by 15.82%. However, *semi-2* yields way better results. 79.94% less error during diastole and 42.87% less error during systole comes to an overall apd of 9.46%. The huge difference for diastolic time steps might be caused by the fact that there are time steps included with a coaptation area of almost zero. Small errors then result in a high apd value. Anyway, it should be mentioned that, to our knowledge, the

coaptation area is not presented for any other (semi-)automatic modelling approach, which might underline the difficulty of delivering a satisfying result for this magnitude.

Furthermore, the plot for an individual data set in figure 79 shows that, while there are significant differences between ground truth and the generated model, a common trend can be seen. So the basic approach to detect the coaptation area is promising, but needs further refinement.

Leaflet areas can be pretty useful in clinical context, for example when comparing MVs before and after a resection, or to determine how much a leaflet stretches during a cardiac cycle. The automatic approach predicted both leaflet areas with 16.3% error, while the anterior leaflet error was slightly smaller. When comparing the mean errors of both, the difference is not as big. This is because the posterior leaflet is smaller, so the same total error has more effect on the posterior leaflet than on the anterior one. Also, the average error during diastole was again significantly better than during systole. Once again, this has to be interpreted as an effect of the coaptation segmentation, which works better for open valves. Using *semi-1*, the error drops by 8.22% on average, while the anterior leaflet difference did merely change during systole and improved by 24.17% during diastole. This different improvement might also be caused by the uneven leaflet size.

Semi-2 again, was able to improve the results a lot. The average apd is at 2.06%, which is an improvement of 87.33%. The systole-specific improvement is a bit higher and another hint that the systolic results for the *auto* segmentation can be affiliated with the difficulty to find the coaptation line in these time steps.

Since the commissure points contribute to the orifice area and orifice circumference, the *semi-1* configuration was able to improve these values by 7.91% and 28.68%.

To sum it up, the *semi-1* configuration does yield better results than *auto*, but the difference is not significant. Surprisingly, the overall improvement was much better during systole. This was unexpected, because the commissure points were detected with less error during diastole. So one would assume that user-defined commissure points would improve systolic time steps more significantly. A conclusion could be that the commissure points contribute far less to the overall segmentation result than assumed. Instead, they improve the coaptation segmentation, which then affects the other values, like coaptation and leaflet area, as well as the leaflet precision D_s . The coaptation segmentation is then again more precise during diastole and this ends up in the fact that *semi-1* improves open valve segmentation more than closed valve segmentation. The *semi-2* configuration was able to deliver good results with respect to the leaflet area. The coaptation area is

with an systolic error of 27.16% not reliable for clinical use.

While the *auto* segmentation predicted the orifice area OA with an average apd of 20.73%, the plotted curve in figure 79 shows that its trend is close to that of the ground truth in the given example. The transition between systole and diastole is also given and suggests that OA could be used to automatically determine the valve state. This may be useful to apply state specific modifications and refinements to a segmented valve.

An evaluation of inter-user-variability would be necessary to decide if the segmentation errors are within the range of user errors. If this is the case, then *semi-2* can deliver segmentations with the quality of a manual segmentation in less time. As Siefert et al. [32] have shown with their simulator, a user error of 0.34 ± 0.30 mm is likely. As a reference point, Graser [6] has examined the intra-user variability for the manual annulus segmentation and found a difference of up to 2.48 ± 0.76 mm for a single user. Generally, the leaflets are easier to detect, but it is hard to tell where the leaflets actually coapt during systole. So the expected intra-user variability for D_s might be smaller, whereas it could be higher for $Coap_x$, for instance.

The automatic segmentation seemed to do well by looking at the model to model and Hausdorff distances. After examining the results in more detail by using the new measurements, it has become clear that the approach can be significantly improved by utilizing user-defined points. This proves that measurements are necessary that are tailored to the given model structure, and semantics if available. Furthermore it casts doubt on the clinical relevance of the results presented in the related work, since general distance measures like D_s and the Hausdorff distance just do not satisfy the requirements.

9.2 Manual Modelling and Ground Truth

Assuming that a 3D+T image should have at least 10 time steps for a cardiac cycle, would mean that a full segmentation takes about a minimum of 80 minutes. Small breaks that are normal during a common work flow are not included, such that a completely manual segmentation is a long and tedious procedure. Subject observation during the evaluation process has shown that the manual segmentation method is not optimal. The double slicing, that was supposed to supply the user with a good choice of slices to cover the valve anatomy without requiring any user interaction, may fail in cases where the coaptation is very close to the annulus. Figure 83 shows a situation where the double slicing fails. The coaptation is very close to the annulus curve and the annulus points corresponding to the red (active) coaptation points are too far off. The resulting slices do not cover the valve

in any reasonable way in this specific example.




Figure 83: Bad double slicing due to a coaptation that is very close to the annulus ring. The slices merely cut the valve.

This example makes clear that the proposed slices should be computed in another way that can ensure a good slicing, or give the user the opportunity to adjust the slices manually. Slices like the ones shown in the example make it hard for the user to delineate the leaflets, since no clear structure is visible in that part of the volume. Consequently it takes longer to model a valve.

The whole evaluation process was done with one expert and all data sets were only segmented once. A more extensive evaluation would include multiple experts segmenting the same data set and the same expert segmenting the same data set multiple times to determine intra- and inter-subject variability of the segmentations.

As implied above and in section 6.1, designing a manual modelling method for the mitral valve is not trivial. This work has put a lot of effort and expert feedback into the development of the proposed manual modelling method, and there are still obstacles to overcome. Thus, the question arises how precise the ground truth models are in the works proposed in section 4, since the model acquisition is not discussed in any thorough detail. Only Pouch refers to [39] as their utilized method and states the time required to segment a single time step to be three to five hours [5].

Schneider [22] uses contours drawn on cut planes by an expert. The contours were then compared to the intersection of each cut plane and the generated model.

Ionasec [21] does not give any information about the manual modelling, except that manual annotations were available for comparison.

The detailed description of the manual segmentation process of the work at hand is one of the advantages of this work, since it makes the evaluation of the (semi-)automatic segmentation quality more comprehensible and reliable.

9.3 Semi-Automatic Segmentation

The evaluation has shown that the segmentation results get better with increasing user input. Anyway, in the field of MV treatment it is intolerable to exchange precision against time. A badly executed surgery can lead to the necessity of another intervention. If a proposed MV segmentation seems to be unreliable it will not be taken into account for surgery planning. So at this point the *semi-2* configuration is the one that has potential to both, save time and deliver precise segmentations. Proposed points for the leaflets have had the most effect, saving at least 43.21% of the manual segmentation time.

The results for the masked and unmasked pipeline were almost equal. A possible explanation for this may be that the coaptation optimization compensates the difference between the two pipeline approaches. The optimization takes only about 1.69 seconds per time step to finish and the masked pipeline takes 65.26 additional seconds per time step to finish and lead to similar results, so a more detailed investigation was neglected in this work. Anyway, future work that focusses on the optimization of the coaptation area prediction might still profit from the masked approach. The leaflet separation must then be further developed to deliver more sophisticated results.

Graser [6] states that an expert can segment the annulus of a volume of 25 time steps in 551 seconds, using the Mitralyzer. This is 22.04 seconds per time step. The sequential unmasked pipeline can process a time step with given annulus in 35.96 seconds. Leaflet segmentation with proposed points takes 177.6 seconds. So the active time that an expert needs to segment a whole mitral model for one time step is at about 200 seconds.

Part IV

Conclusion

10 Summary

This work has been motivated by the fact that the interpretation of 2D US images can be very difficult and that the most relevant decisions during a MVR can only be made as soon as the surgeon has direct vision on the patient's MV. A pre-operative segmentation of the MV can support the surgeon during an intervention.

Chapter II introduces the background of this work. Section 2 describes the anatomy of the MV and the effects of MR, as well as basic information about how a diseased MV can be repaired. The technical background is outlined in section 3. The TEE imaging modality is leading when it comes to 3D image quality, but has the drawback that it can only be applied to anaesthetized patients, since the ultrasound probe has to be inserted into the patient's esophagus. Further, segmentation methods that were utilized in this work are presented in this section.

Section 4 introduces the MITK, that was used as the base framework for the implementation of this work. Methods of other research groups that work in the field MV segmentation are also proposed in this section. The investigation of the related work has shown that various approaches seem to yield results of similar quality, but comparability of different works is hardly given, since each method uses their own mitral model representation. Furthermore, the utilized measurements to obtain the segmentation quality have not proven to be reliable. Additionally, two commercial work stations that offer MV segmentation methods are pictured.

Chapter III contains the main content of the work at hand. Section 5 states further benefits that can be taken from a segmented valve for surgeons and researchers. Then the final objectives for the manual and (semi-)automatic segmentation methods are defined.

The following section deals with the theory and concepts that form the basics of the manual and (semi-)automatic methods and introduces the model representation used in this work. Each component of the automatic segmentation pipeline is described in detail.

The experiments and their results are provided in section 8. This section includes the proposal of the new magnitudes, that were used to measure the segmentation quality of generated models, and the description of the evaluation process. The results are presented for the semi-automatic and manual segmentation and for the runtime analysis of the prototypical im-

plementation of the algorithm proposed in the work at hand. A final discussion and interpretation of the results is done in section 9.

11 Contributions

This section will briefly summarize the fulfilled objectives of this work. The following overall goals have been achieved:

- Fully automatic segmentation method for the mitral leaflets, with an annulus as input.
- Manual segmentation method for the mitral leaflets.
- Automatic and manual segmentation can be combined to a user-guided semi-automatic process.
- The result measurements are on the order of related state of the art methods.
- Novel, clinically relevant measurements can be obtained from the segmentations.
- User effort can be drastically reduced to obtain models of manual segmentation quality.

The objectives listed in section 5 have been accomplished, but correspondences between measurement results and the pathologies of different subjects were not found, since the number of used data sets was too small to make a reliable statement. The proposed method is capable of executing a fully automatic segmentation, with an annulus ring as input (1). Parts of the automatic process can be used to speed up the proposed manual segmentation method (2,3,4). The models created by the automatic configuration are on the level of other state of the art methods (5), but thorough evaluation has shown that more user input can significantly increase the quality of a segmentation. Semantic information in the model representation allows for measuring the coaptation area, commissure points and length of the leaflet ends, which can be relevant in clinical use (6).

The objectives for the manual segmentation in section 5.2 have also been completed. An annulus model can be used as input (1) to segment the leaflet tips and surface and the papillary muscles over a whole cardiac cycle (4,5,6,7). The process is divided into several steps (2) and the user can go back and forth between those steps for editing (8). Though objectives 7 and 8 are implemented, they require further improvement. Information of multiple time steps could be displayed to the user for guidance during

the manual segmentation and the single steps are recommended to be executed in a certain order to achieve the best efficiency.

The semi-automatic objectives were fulfilled as well. The method is based on a given annulus ring (1). Leaflets can be automatically separated without additional user-input (2), mainly based on and inspired by Schneider's graph cut method (3). New methods have been proposed to fit a 3D model into pre-processed image data (4) and the manual segmentation method can be used to guide the automatic pipeline steps.

As a recapitulation, this work has introduced a new segmentation method for the mitral leaflets. This method is partially based on other existing work but has contributed new ideas for leaflet separation and the integration into a manual segmentation work-flow. It was found that even a manual approach has to be planned thoroughly, to achieve an efficiently usable modelling environment, that proposes the necessary visual information to the user and results in anatomically plausible models. The combination of the automatic and manual modelling approach has resulted in an opportunity to obtain segmentations of manually modelled quality, while drastically reducing labour time. Furthermore the work at hand has revealed that commonly used methods to measure the segmentation quality of MVs are not satisfying. A more detailed inspection of anatomical landmarks of the result models, using new measurement methods, has shown that it is very difficult to generate models that can be used for clinically relevant measurements.

12 Future Work

The manual as well as the automatic segmentation in this work are facing several obstacles and difficulties that restrict the portion of data sets that can be processed with the proposed methods or impair the result quality. Therefore future work could deal with the improvement of the existing methods to make them more general and sophisticated, instead of reaching for new ideas.

Anyway, there is one general goal that could improve the research in the field of MV segmentation and that would be a general model representation. The works mentioned in section 4 and this work show, that practically each research group uses their own model representation. Comparing segmentation quality among different groups is very difficult because of that. Also, it is hard or even impossible to combine several methods, if they are tailored to a specific representation. A globally recognized model concept of the MV would surely contribute a lot to the research in this field. A

model of this kind should be compatible with simulation, visualization, efficient access to the data for segmentation tasks, anatomical plausibility with medical semantics and intuitive modelling concepts for manual segmentation.

The following sections describe a couple of promising thoughts that apply to the work at hand.

12.1 Model Representation

The representation for the MV model should be further developed to better fit the actual anatomy of the MV. The model used in this work uses the same number of subdivisions for the anterior and posterior leaflet although the posterior leaflet has a much smaller area. Consequently the information supplied for the posterior leaflet is much more dense than for the anterior leaflet. Pouch et al. have shown that the leaflet thickness can be detected and used to determine the healthiness of a MV. This information is completely missing in this work. Furthermore the model should be more flexible. The whole segmentation process is linked to the few landmarks that delineate the model and internally each coaptation point has a fixed link to an annulus point. Because of the various shapes that a MV can take on, these links can be very close to each other, resulting in very dense information, or very far off, which neglects large portions of the leaflet anatomy. A redesign of the model representation should guarantee an equal spatial sampling.

12.2 Shape Prior

The segmentation algorithm described in this work and the manual segmentation approach make it possible for an expert to create segmentations of a whole cardiac cycle with less effort. A wide collection of segmentations could be built up and examined with this. As mentioned in section 3.2.1, the number of data sets available for this work were found to not suffice to generate a shape prior that can describe the various shapes of the MV. A future refinement of the model initialization and optimization process could be augmented by a MV shape prior whose creation is possible with the less labour-intensive semi-automatic modelling. The SP might contribute to a better prediction of anatomical landmarks and maintenance of a time consistent segmentation. This would be another advantage of a general model representation. The models gathered by individual researchers could be combined to a SP.

Figure 83 has shown that the double slicing method does not present satisfying slices in all cases. So, a step to be done before creating a SP would be

to redesign the slicing method for manual leaflet modelling. The fan slicing has proven to be well suited to slice the valve area between the antero-lateral and postero-medial commissure points, with the drawback that the remaining area is not covered so well (see figure 32). Figure 84 shows a suggestion for a new slicing method. The center point that the fan slices are rotated about is set further off the annulus, above the saddle horn. This ensures that the area covered on the anterior and posterior leaflet is more equal. The remaining area is then covered with slices centered at each commissure point. This should assure that the whole valve can be modelled and there is no relation between coaptation and annulus points that can lead to impractical slices.

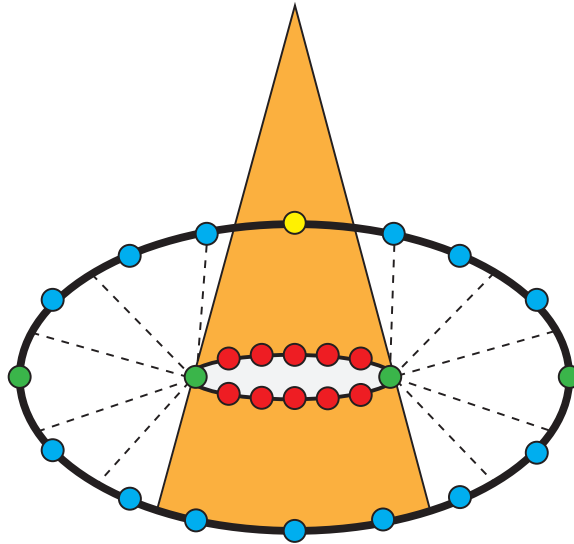


Figure 84: Suggestion for a new slicing method.

12.3 Inter- and Intra-user Variability

The evaluation for the new measurements proposed in this work was done by comparing the (semi)-automatically generated models against the user-defined ground truth, so it could be shown how much the results differ from the segmentation of that one expert. For further evaluation the generated models should also be compared against the segmentations of other experts. Multiple segmentations of the same data set could be averaged to a mean ground truth of that data set. Also, it should be determined how much variance can be found among experts and among multiple segmentations of the same data set and the same expert. This would reveal how close the automatic algorithm is to a manual segmentation of an unknown expert.

Furthermore an examination like this could provide information on how

difficult it is to model a certain region of the MV. For instance, a small variance for the coaptation lines would mean that this part is easy to detect by a human and a reliable source of information. Whereas, if the commissure points vary a lot among different users and segmentations, this could be a hint that there is not "that" commissure point, but rather a commissure region - or that the commissure point is simply very hard to find. Landmarks that can be modelled precisely should be taken more into account for surgical planning than landmarks that can not be certainly pointed at.

12.4 3D Leaflet Separation

The proposed segmentation algorithm in this work can not cope with pathologies where the leaflets overlap, because of the 2D based leaflet separation and MIP. Future development could try to extend the 2D graph cut method to a 3D method, to overcome this restriction. Figure 85 shows the drawback of the 2D graph cut. If the leaflets overlap, but do not touch in 3D, the information about the gap is lost after the MIP, which joins the leaflets.




Figure 85: Slice of a 3D grid. Gray squares show voxels classified as tissue. The MIP of the slice is shown at the bottom. The 3D graph cut is shown as a dashed line.

The nodes are connected to sink and source as described in section 6.3.6, but the topology is extended to 3D. After the 3D separation, the model initialization step could be done individually on a MIP of each leaflet. Utilizing 3D information for separation of the leaflets is a crucial step to make the segmentation method of this work more general.

12.5 SeMitralyzer Work-flow

The prototypical implementation of the SeMitralyzer in this work does not allow for an efficient work-flow. At present, the annulus has to be modelled for all time steps, then all time steps are segmented at once (parallelized or sequential). After that, the user can go on to model the commissure, coaptation and leaflets with proposed points.

An efficient work-flow might be designed as follows: The user could at first segment all annulus curves for all time steps and for each finished annulus time step, the leaflet segmentation for that time step starts executing in the background. While the SeMitralyzer processes a time step in 35.95 seconds and a manual annulus segmentation takes 22.04 seconds, the leaflet segmentation for at least one time step should have finished before the user has manually segmented the annulus for all time steps. The leaflet segmentation then takes 177.6 seconds with proposed points. In the meantime the leaflet segmentation for the next time step is definitely finished.

This way, there would be no waiting time for a user during the semi-automatic segmentation.

References

- [1] American Heart Association. Online at:
http://www.heart.org/HEARTORG/Conditions/More/HeartValveProblemsandDisease/Problem-Mitral-Valve-Prolapse_UCM_450441_Article.jsp
accessed July 2014.
- [2] U.S News and World Report. Online at:
<http://www.reparacionvalvularmitral.org/images/press/usnews.pdf>
accessed July 2014.
- [3] Universitätsherzzentrum Freiburg. Online at:
<http://www.herzzentrum.de/UHZ/pdf/HZA/UHZ-aktuell-Jan13pdf.pdf>
accessed July 2014.
- [4] Roberto M. Lang, Wendy Tsang, Lynn Weinert, Victor Mor-Avi, and Sonal Chandra. Valvular heart disease. 58(19):1933–1944.
- [5] Alison M. Pouch, Paul A. Yushkevich, Benjamin M. Jackson, Arminder S. Jassar, Mathieu Vergnat, Joseph H. Gorman, Robert C. Gorman, and Chandra M. Sehgal. Development of a semi-automated method for mitral valve modeling with medial axis representation using 3d ultrasound. 39(2):933–950.
- [6] Bastian Graser. Computer assisted annuloplasty planning.
- [7] Z. G. Turi. Mitral valve disease. 109(6):38e–41.
- [8] Alain Carpentier, David H Adams, and Farzan Filsoufi. *Carpentier's reconstructive valve surgery from valve analysis to valve reconstruction*. Saunders/Elsevier.
- [9] John E. Aldrich. Basic physics of ultrasound imaging:. 35:S131–S137.
- [10] Robert J. Schneider, Douglas P. Perrin, Nikolay V. Vasilyev, Gerald R. Marx, Pedro J. del Nido, and Robert D. Howe. Mitral annulus segmentation from 3d ultrasound using graph cuts. 29(9):1676–1687.
- [11] Judy Hung, Roberto Lang, Frank Flachskampf, Stanton K. Shernan, Marti L. McCulloch, David B. Adams, James Thomas, Mani Vannan, and Thomas Ryan. 3d echocardiography: A review of the current status and future directions. 20(3):213–233.

- [12] A.M. Pouch, H. Wang, M. Takabe, B.M. Jackson, J.H. Gorman, R.C. Gorman, P.A. Yushkevich, and C.M. Sehgal. Fully automatic segmentation of the mitral leaflets in 3d transesophageal echocardiographic images using multi-atlas joint label fusion and deformable medial modeling. 18(1):118–129.
- [13] Tim McInerney and Demetri Terzopoulos. Deformable models in medical image analysis. In *Mathematical Methods in Biomedical Image Analysis, 1996., Proceedings of the Workshop on*, pages 171–180. IEEE.
- [14] Tim Cootes. Model-based methods in analysis of biomedical images. pages 223–47.
- [15] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *INTERNATIONAL JOURNAL OF COMPUTER VISION*, 1(4):321–331, 1988.
- [16] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models—their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59, January 1995.
- [17] Y. Boykov and V. Kolmogorov. An experimental comparison of mincut/max-flow algorithms for energy minimization in vision. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 26, pages 1124–1137.
- [18] Marco Nolden, Sascha Zelzer, Alexander Seitel, Diana Wald, Michael Müller, Alfred M. Franz, Daniel Maleike, Markus Fangerau, Matthias Baumhauer, Lena Maier-Hein, Klaus H. Maier-Hein, Hans Peter Meinzer, and Ivo Wolf. The medical imaging interaction toolkit: challenges and advances: 10 years of open-source development. 8(4):607–620.
- [19] Mathias Seitel. Bildbasierte planungsunterstützung chirurgischer eingriffe bei erkrankungen der mitralklappe.
- [20] Razvan Ioan Ionasec, Bogdan Georgescu, Eva Gassner, Sebastian Vogt, Oliver Kutter, Michael Scheuering, Nassir Navab, and Dorin Comaniciu. Dynamic model-driven quantitative and visual evaluation of the aortic valve from 4d CT. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2008*, pages 686–694. Springer.
- [21] Razvan Ioan Ionasec, Ingmar Voigt, Bogdan Georgescu, Yang Wang, Helene Houle, Fernando Vega-Higuera, Nassir Navab, and Dorin Comaniciu. Patient-specific modeling and quantification of the aortic and mitral valves from 4-d cardiac CT and TEE. 29(9):1636–1651.

- [22] Robert J. Schneider, William C. Burke, Gerald R. Marx, J. Pedro, and Robert D. Howe. Modeling mitral valve leaflets from three-dimensional ultrasound. In *Functional Imaging and Modeling of the Heart*, pages 215–222. Springer.
- [23] Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–511. IEEE.
- [24] William T. Freeman Edward H. Adelson. The design and use of steerable filters.
- [25] Zhuowen Tu. Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering. In *Proceedings of the Tenth IEEE International Conference on Computer Vision - Volume 2, ICCV '05*, pages 1589–1596, Washington, DC, USA, 2005. IEEE Computer Society.
- [26] Robert J. Schneider, Douglas P. Perrin, Nikolay V. Vasilyev, Gerald R. Marx, Pedro J. del Nido, and Robert D. Howe. Mitral annulus segmentation from four-dimensional ultrasound using a valve state predictor and constrained optical flow. *16(2):497–504*.
- [27] Robert J. Schneider, Neil A. Tenenholtz, Douglas P. Perrin, Gerald R. Marx, J. Pedro, and Robert D. Howe. Patient-specific mitral leaflet segmentation from 4d ultrasound. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2011*, pages 520–527. Springer.
- [28] Adam Huang, Gregory M. Nielson, Anshuman Razdan, Gerald E. Farin, D. Page Baluch, and David G. Capco. Thin structure segmentation and visualization in three-dimensional biomedical images: A shape-based approach. *IEEE Trans. Vis. Comput. Graph.*, *12(1):93–102*, 2006.
- [29] Sébastien Martin, Vincent Daanen, Olivier Chavanon, and Jocelyne Troccaz. *Fast segmentation of the mitral valve leaflet in echocardiography*. Springer.
- [30] Philippe Burlina, Chad Sprouse, Ryan Mukherjee, Daniel DeMenthon, and Theodore Abraham. Patient-specific mitral valve closure prediction using 3d echocardiography. *39(5):769–783*.
- [31] Ingmar Voigt, Tommaso Mansi, Razvan Ioan Ionasec, Etienne Assoumou Mengue, Helene Houle, Bogdan Georgescu, Joachim Hornegger, and Dorin Comaniciu. Robust physically-constrained modeling

- of the mitral valve and subvalvular apparatus. In Gabor Fichtinger, Anne L. Martel, and Terry M. Peters, editors, *MICCAI (3)*, volume 6893 of *Lecture Notes in Computer Science*, pages 504–511. Springer, 2011.
- [32] Andrew W. Siefert, David A. Icenogle, Jean-Pierre M. Rabbah, Neelakantan Saikrishnan, Jarek Rossignac, Stamatios Lerakis, and Ajit P. Yoganathan. Accuracy of a mitral valve segmentation method using j-splines for real-time 3d echocardiography data. 41(6):1258–1268.
- [33] QLAB - Philipps. Online at:
www.healthcare.philips.com
 accessed July 2014.
- [34] TomTec. Online at:
www.tomtec.de
 accessed July 2014.
- [35] Haider Javed Warraich, Sajid Shahul, Robina Matyal, and Feroze Mahmood. Bench to bedside: dynamic mitral valve assessment. *Journal of cardiothoracic and vascular anesthesia*, 25(5):863–866, October 2011.
- [36] Felix Hausdorff. *Grundzüge der Mengenlehre*. Veit and Company, Leipzig, 1914. Das Hauptwerk von Felix Hausdorff.
- [37] Günter Rote. Computing the minimum hausdorff distance between two point sets on a line under translation. *Inf. Process. Lett.*, 38(3):123–127, May 1991.
- [38] Jonathan S. Grashow, Ajit P. Yoganathan, and Michael S. Sacks. Biaxial stress-stretch behavior of the mitral valve anterior leaflet at physiologic strain rates. *Annals of Biomedical Engineering*, 34(2):315–325, 2006.
- [39] A. S. Jassar, C. J. Brinster, M. Vergnat, J. D. Robb, T. J. Eperjesi, A. M. Pouch, A. T. Cheung, S. J. Weiss, M. A. Acker, J. H. 3rd Gorman, R. C. Gorman, and B. M. Jackson. Quantitative mitral valve modeling using real-time three-dimensional echocardiography: technique and repeatability. *Ann Thorac Surg*, 91:165–71, 2011. 1.