CHALMERS
UNIVERSITY OF TECHNOLOGY

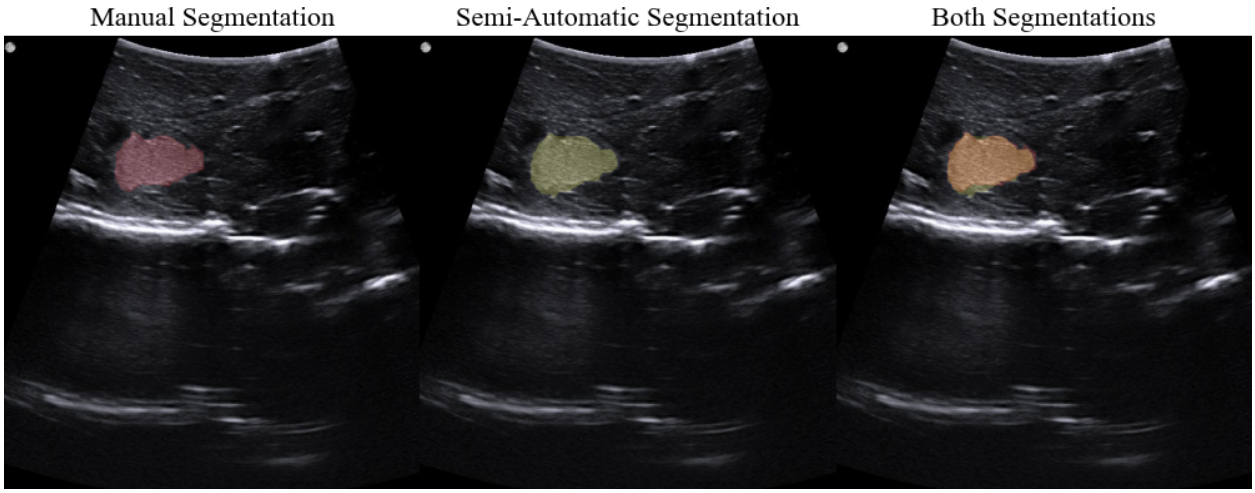

Semi-Automatic Segmentation via Ultrasound Imaging

For AR-Guided Laparoscopic Liver Tumor Resections

Master's thesis in Biomedical Engineering

ALBIN WIDENGÅRD

DEPARTMENT OF MATHEMATICAL SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2024
www.chalmers.se

Master's thesis 2024

Semi-Automatic Segmentation via Ultrasound Imaging

For AR-Guided Laparoscopic Liver Tumor Resections

ALBIN WIDENGÅRD

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mathematical Sciences
Chalmers University of Technology
Gothenburg, Sweden 2024

Semi-Automatic Segmentation via Ultrasound Imaging
For AR-Guided Laparoscopic Liver Tumor Resections
ALBIN WIDENGÅRD

Supervisor: Klas Modin, Professor, Mathematical Sciences
Examiner: Klas Modin, Professor, Mathematical Sciences

Master's Thesis 2024
Department of Mathematical Sciences
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: The developed semi-automatic segmentation method compared to a manual segmentation.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2024

Semi-Automatic Segmentation via Ultrasound Imaging
For AR-Guided Laparoscopic Liver Tumor Resections
ALBIN WIDENGÅRD
Department of Mathematical Sciences
Chalmers University of Technology

Abstract

Liver cancer is one of the most prevalent and deadly forms of cancer. The primary treatment for liver cancer is the surgical removal of the tumor, known as resection. While traditionally performed via open surgery, there has been a significant shift towards minimally invasive surgery, which involves smaller incisions and offers advantages such as faster recovery and shorter hospital stays. Despite its benefits, minimally invasive surgery presents challenges in distinguishing between tumor tissue and healthy liver tissue. An effective solution to these challenges is an augmented reality guiding tool that assists surgeons by overlaying a 3D model of the tumor onto the laparoscopic camera view during surgery.

This master's thesis, conducted in collaboration with Navari Surgical, aims to develop a semi-automatic segmentation method based on ultrasound imaging for creating such a 3D tumor model. The method involves preprocessing techniques and a segmentation algorithm implemented in Python. The evaluation, based on relevant metrics, demonstrates that the method performs well with images of high tumor visibility, though challenges remain with images containing multiple tumor areas and those requiring individualized preprocessing parameters. Future work should focus on refining these aspects and incorporating machine learning models to enhance accuracy and usability. This study establishes the feasibility of using ultrasound imaging for liver tumor segmentation and provides a foundation for further research in this area.

Keywords: segmentation, ultrasound, liver tumor, laparoscopy, augmented reality, active contour models.

Acknowledgements

I would like to express my heartfelt appreciation to the amazing team at Navari. Your camaraderie and positivity have made this journey incredibly enjoyable. A special thanks to Carl Bodin for his continuous support in discussing this master's thesis. From the smallest details to the overall layout, your guidance has been invaluable. Thank you also to Madeleine Gustavsson and Axel Blomé for your unwavering support and belief in me from the very beginning.

Additionally, I extend my deepest gratitude to my examiner, Klas Modin, for his insightful feedback and constructive criticism throughout this process. Your support has been instrumental in the completion of this thesis.

Albin Widengård, Gothenburg, June 2024

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

CT	Computerized Tomography
GUI	Graphical User Interface
MIS	Minimally Invasive Surgery
ML	Machine Learning
MR	Magnetic Resonance
SAM	Segment Anything Model
US	Ultrasound

Contents

List of Figures

List of Tables

1
Introduction

This section provides a comprehensive introduction to the topic, detailing the overall aim and identifying key limitations of the project.

1.1 Background

The primary treatment for liver cancer is the surgical removal of the tumor, known as resection [3]. Traditionally, this procedure was performed via open surgery, requiring a large incision in the abdomen. However, there has been a significant shift towards minimally invasive surgery (MIS), which involves making small incisions in the abdomen through which a laparoscopic camera and various surgical tools can be inserted. This approach offers numerous advantages including faster patient recovery, shorter hospital stays and the avoidance of large abdominal scars. Laparoscopy, commonly referred to as keyhole surgery, is a prevalent form of MIS.

Despite its benefits, MIS has inherent challenges. During laparoscopic liver tumor resections, surgeons often struggle to distinguish between tumor tissue and healthy liver tissue [4]. Surgeons must mentally visualize the tumor contours based on pre-operative CT and MR scans, which demands high cognitive effort and is time-consuming. Additionally, surgeons frequently use laparoscopic ultrasound (US) intraoperatively to verify their progress, further adding to the procedure's duration.

An effective solution to these challenges is an augmented reality (AR) guiding tool that assists surgeons during operations. This master's thesis is conducted in collaboration with Navari Surgical, a medical technology startup developing such a solution. Their product enhances visual feedback by overlaying a 3D model of the tumor onto the laparoscopic camera view used during surgery, significantly reducing cognitive effort and improving accuracy.

Traditionally, creating this 3D model relies on CT imaging, which necessitates hybrid operating rooms equipped with CT machines. This requirement poses a limitation as hybrid operating rooms are not always accessible and do not fit seamlessly into existing surgical workflows. In contrast, using US to create these models presents several advantages: it is less expensive, more accessible, does not expose the patient to ionizing radiation and does not require a hybrid operating room.

A crucial step in generating the 3D tumor model from US images is the segmenta-

tion process, which involves identifying and outlining the tumor boundaries in each 2D US image. This master's thesis aims to develop a semi-automatic segmentation method to facilitate this process.

1.2 Aim

The aim of this master's thesis is to develop a semi-automatic segmentation method based on ultrasound imaging. The method aims to be used as part of the creation of a 3D tumor model for an AR guiding tool during laparoscopic liver tumor resections by Navari Surgical. This involves exploring appropriate preprocessing techniques, implementing the segmentation method in Python and evaluating its effectiveness using relevant metrics. The method's suitability, performance and potential will then be determined. Additionally, this thesis will serve as a feasibility study to determine the viability of using ultrasound imaging for liver tumor segmentation.

1.3 Limitations

Due to the limited timeframe of this master's thesis project, several constraints were necessary to ensure its success. Firstly, the scope of the thesis is restricted to 2D segmentation and does not encompass the reconstruction of a 3D tumor model or segmentation in 3D. Additionally, limitations exist both within the dataset and the developed segmentation method itself. The images used for the development of the method are relatively straightforward to segment and there are certain constraints related to the preprocessing steps, which affect the method's overall readiness. A comprehensive discussion of all limitations can be found in Section 5.2.

1.4 Prior Work

There are research groups that have managed to create laparoscopic 3D reconstructions with US. One of them is a group from London, Canada, that in 2017 successfully managed to visualize a tumor phantom [4]. The method used was to stitch together 2D US images into a 3D volume using a high-performance reconstruction algorithm and then display that volume within the laparoscopic image. A study involving laparoscopic surgeons showed that the cognitive and physical effort in visualizing the hidden tumor is significantly lower using this method than the conventional one. However, at greater depths, the conventional US method showed to be more accurate in depth judgment. During the author's conversation with this research group in 2023, they were extremely positive to the feasibility of US being used for AR-guided laparoscopic surgeries.

There are also studies on semi-automatic segmentation methods of liver tumors. A greek research group developed a semi-automatic and hybrid semi-automatic liver segmentation algorithm in 2015 [5]. In the semi-automatic method, the user selects initial seed points (typically five) on a slice of the medical image. The algorithm

performs segmentation based on thresholding, using the mean or median pixel intensity value from neighboring pixels around the seed points. This process involves morphological and binary image filter operations, such as spur pixel deletion and average binary filtering, to produce an intermediate binary segmentation image. The final segmentation mask is generated from binary objects overlapping the initial seed points. The hybrid semi-automatic algorithm processes the entire volume dataset at once. Initially, the user selects seed points on a desired slice, similar to the semi-automatic method. The algorithm then automatically determines seed points for the remaining slices (default is 15 automatic seed points per slice). This sequential processing starts from the user-processed slice and moves downwards and upwards through the dataset. The seed points for each subsequent slice are derived from the segmentation mask of the previous slice. Segmentation on each slice follows the same semi-automatic algorithm with user-defined threshold limits. The results of this study were good, however, the segmentation relied on CT and MR scans and not US.

This master's thesis is the first work done at Navari involving semi-automatic segmentation of US images.

2
Theory

In this chapter, the necessary theory for the master's thesis project will be presented. Initially, information regarding the liver and liver cancer is discussed, followed by an overview of ultrasound imaging. Subsequently, the principles of image processing and segmentation methods, including classical methods and active contour models, are explored. The chapter also delves into specific techniques such as Morphological Snakes, which is integral to the project's methodology. Finally, the evaluation metrics used to assess the developed segmentation method are introduced.

2.1 Medical Background

In this section, medical information about the liver and liver cancer will be presented.

2.1.1 Liver Anatomy

The liver, the largest internal organ in the human body, performs several essential and complex functions [6]. Key functions include maintaining blood glucose levels through glycogen storage, synthesizing proteins, detoxifying various substances and producing bile. Positioned beneath the diaphragm in the upper right quadrant of the abdomen, the liver typically weighs between 1300 to 1700 grams in adults and possesses a spongy, continuous texture.

Predominantly, the liver comprises two unique cell types: hepatocytes and cholangiocytes [6]. Hepatocytes are the predominant cell type, tasked with critical functions such as storing glycogen, detoxifying harmful substances and bile production. Cholangiocytes are crucial for bile drainage; they construct the biliary channels and alter bile composition through the secretion and absorption of various compounds.

The liver also features a distinctive vascular architecture, receiving blood from two primary sources: approximately 75% from the portal vein and 25% from the hepatic artery [6]. The portal vein transports nutrient-rich, deoxygenated blood from the gastrointestinal tract, allowing the liver to process substances before they enter the systemic circulation. Conversely, the hepatic artery delivers oxygenated blood from the systemic circulation to the liver. These vessels subdivide into arterioles and venules, which discharge into a capillary network of specialized capillaries known as sinusoids. Blood from the sinusoids collects in small veins that exit the liver through the hepatic veins into the inferior vena cava, integrating into the systemic

circulation. Roughly 22% of the liver's mass or volume comprises blood vessels, designed to accommodate substantial blood volumes at high flow rates; about 12% of the body's total blood volume is held within the liver under physiological conditions.

Functionally, the liver is segmented into eight distinct regions based on vascular supply, with each segment being an autonomous unit with its own vascular and biliary branches [6, 3]. This segmentation facilitates the surgical resection of one or several segments without impacting the remaining liver tissue. Liver cells, including hepatocytes and cholangiocytes, can proliferate to regenerate tissue, mirroring the characteristics of undamaged liver tissue. Remarkably, up to 70% of the liver can regenerate, a vital feature after surgical resections, allowing the remaining segments to enlarge and compensate for the lost tissue.

2.1.2 Liver Cancer

Liver cancer is categorized into primary and secondary types. Primary liver cancer begins in the liver, whereas secondary liver cancer starts elsewhere in the body and spreads to the liver. Primary liver cancer ranks as the fifth most common cancer globally and is the second leading cause of cancer-related deaths [7]. The two main forms of primary liver cancer, hepatocellular carcinoma (HCC) and intrahepatic cholangiocarcinoma (ICC), constitute over 95% of all primary liver cancer cases. HCC, the predominant type, accounts for about 80-85% of these cases and typically develops from hepatocytes. It frequently occurs in individuals with chronic liver conditions such as hepatitis B and C. In contrast, ICC originates from the bile ducts within the liver [8].

Treatment strategies for HCC are highly personalized, considering factors like the stage of the disease, the extent of liver dysfunction and the patient's overall health and age [7, 6]. For individuals without cirrhosis and who maintain good liver function, tumor resection is often the preferred treatment. The liver's regenerative capacity enables the resection of even large tumors. It is crucial to remove all tumor tissue during resection to ensure the functionality of the remaining liver tissue is not compromised. However, resection may not be advisable for patients with severe chronic liver disease or cirrhosis due to the heightened risk of postoperative complications. Patients with metastatic liver cancer are also not suitable candidates for resection.

For those with advanced liver cirrhosis, liver transplantation is considered the optimal treatment as it can also address the underlying chronic liver condition [7]. Nonetheless, this approach is challenged by the scarcity of donor organs and the complexity of the transplant procedure. An alternative is percutaneous ablation techniques such as radiofrequency ablation, where a needle emitting an electrical current is used to heat and destroy tumor cells [6]. This method is an option for patients ineligible for resection and, in some instances, can yield results akin to those of surgical resection. Other ablation techniques and various pharmacological treatments are also available.

2.2 Ultrasound Imaging

This section presents the theory behind ultrasound imaging, including the underlying physics and potential artifacts that may occur during the imaging process.

2.2.1 Ultrasound Physics

Ultrasound technology employs sound waves with frequencies above the human auditory threshold, generally starting at 20 kHz [9]. These mechanical waves propagate through media such as fluids and solids by causing microscopic particle displacements, facilitating energy transfer without actual mass movement. The behavior of these waves is governed by the three-dimensional, partial differential wave equation

$$\nabla^2 p - \frac{1}{c^2} \frac{\partial^2 p}{\partial t^2} = 0, \tag{2.1}$$

where p is the acoustic pressure, c is the speed of sound, ∇ denotes the Laplacian operator and t represents time [10].

In an ideal, non-viscous fluid, ultrasound waves manifest as longitudinal waves, with particles oscillating in the direction of wave travel [9]. The propagation and interaction of ultrasound waves with different media depend significantly on the medium's specific acoustic impedance Z, defined as

$$Z = \rho c, \tag{2.2}$$

where ρ is the density of the medium. This property is crucial as it affects how ultrasound waves are reflected or transmitted at material interfaces. The energy carried by ultrasound waves can be expressed in terms of acoustic intensity

$$I = \frac{p^2}{2Z}, \tag{2.3}$$

measured in $\mathrm{W/cm^2}$.

2.2.2 Imaging Artifacts

Sonographic artifacts are prevalent in US imaging and must be considered [11]. These can distort the size, position, and shape of structures or even display structures that do not exist. Common artifacts include shadow artifacts, which occur when ultrasound waves encounter solid structures like stones or ribs, creating shadows behind them, useful for diagnosing conditions like gallstones. Posterior enhancement happens when imaging fluid-filled structures such as the gall bladder or urinary bladder, allowing more ultrasound waves to pass through and creating a bright area behind the structure. This artifact requires careful adjustment of gain settings to avoid missing small amounts of fluid. Edge artifacts occur when the ultrasound beam bends at the edge of rounded structures, such as the kidney or urinary bladder, and can be clarified by changing the angle of the beam. Mirror

artifacts appear when ultrasound waves reflect off high acoustic impedance tissues like the diaphragm, producing a virtual, hypoechoic, and slightly blurred image on the opposite side of the original structure. Reverberation artifacts are formed when ultrasound waves bounce between two interfaces with high acoustic impedance, such as the pleura, resulting in parallel lines with equal spacing and decreased density at greater depths, creating a striped pattern.

2.3 Image Processing

Theory behind image processing techniques used in this project are explained in this section.

2.3.1 OpenCV

OpenCV (Open Source Computer Vision Library), initially developed by Intel, is a tool in computer vision, featuring an extensive suite of over 2,500 optimized algorithms [12]. This library facilitates a broad range of tasks from simple image manipulation to complex real-time image analysis, making it essential for both academic and commercial applications.

OpenCV's robust functionality includes operations such as pixel access and matrix manipulation which are critical for filtering, edge detection and geometric transformations. These capabilities are essential for applications that require high-performance image analysis, such as feature detection and automated visual inspections.

2.3.2 Active Contour Models

Active contour models, commonly known as "snakes," are sophisticated tools in image processing designed to identify object boundaries [13]. By iteratively adjusting a spline, these models minimize an energy function to accurately fit to the object's contours.

The total energy of the snake, E^*_{snake}, is described by the integral of internal, image and external energies along the contour:

$$E^*_{\text{snake}} = \int_0^1 \Big(E_{\text{internal}}(\mathbf{v}(s)) + E_{\text{image}}(\mathbf{v}(s)) + E_{\text{con}}(\mathbf{v}(s)) \Big)\, ds, \tag{2.4}$$

where s parameterizes the contour through $\mathbf{v}(s) = (x(s), y(s))$. The internal energy, E_{internal}, acts to stabilize the spline by penalizing excessive bending and stretching and is given by:

$$E_{\text{internal}} = \frac{1}{2} \left(\alpha(s) |\mathbf{v}_s(s)|^2 + \beta(s) |\mathbf{v}_{ss}(s)|^2 \right), \tag{2.5}$$

where $\alpha(s)$ and $\beta(s)$ are parameters that control the spline's tension and rigidity. The image energy, E_{image}, helps to drive the snake towards important features such

as lines and edges. It is formulated as

$$E_{\text{image}} = w_{\text{line}}E_{\text{line}} + w_{\text{edge}}E_{\text{edge}} + w_{\text{term}}E_{\text{term}}, \tag{2.6}$$

with each component attracting the snake to different image characteristics. $E_{\text{line}} = I(x, y)$ targets areas of particular light or dark intensity based on the line weight w_{line}. $E_{\text{edge}} = -|\nabla I(x, y)|^2$ focuses on areas with significant gradients, indicating edges. The termination energy, E_{term}, detects corners or endpoints in the image and is expressed as

$$E_{\text{term}} = \frac{\partial^2 C}{\partial n_\perp^2}\frac{\partial C}{\partial n_\perp} = \frac{C_{yy}C_x^2 - 2C_{xy}C_xC_y + C_{xx}C_y^2}{(C_x^2 + C_y^2)^{3/2}}, \tag{2.7}$$

where $C(x, y) = G_\sigma(x, y) * I(x, y)$ represents the image smoothed with a Gaussian filter, C_{xy}, C_{xx} and C_{yy} are second-order partial derivatives and C_x and C_y are the first-order derivatives in the x and y directions, respectively. Lastly, external constraint energy, E_{con}, is applied to direct the snake towards the desired object or away from undesired local minima. This energy allows for additional user control over the snake's behavior, ensuring that it conforms closely to the target object's contours.

2.3.3 Morphology

Morphological operations are a set of image processing techniques that manipulate images based on their shapes [1]. These operations involve the application of a structuring element to an input image, resulting in the generation of an output image.The cornerstone morphological operations include erosion and dilation, each serving diverse purposes such as noise reduction, element isolation, and feature detection.

Dilation entails convolving an image A with a kernel B, typically of square or circular shape [1]. The kernel B features an anchor point, often placed at its center. As the kernel traverses the image, the maximal pixel value overlapped by B is computed and used to replace the pixel at the anchor point. This process, known as maximizing, causes bright regions within the image to expand, thus the term dilation. See Figure 2.1 for visual representation.

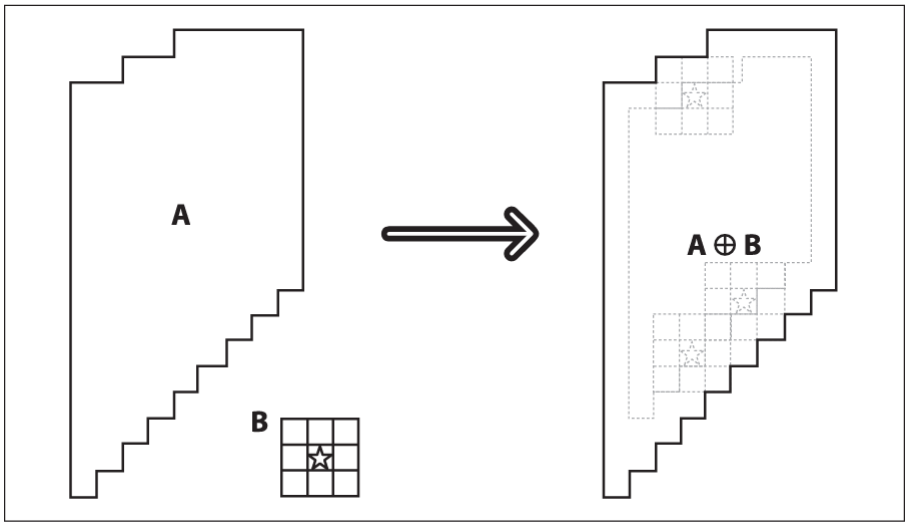

Figure 2.1: Visual representation of the morphological dilation process. Image source: [1].

Contrary to dilation, erosion computes a local minimum over the area defined by the kernel [1]. Similar to dilation, as the kernel B traverses the image, the minimal pixel value overlapped by B is determined and used to replace the pixel under the anchor point. This process results in the thinning of bright areas and the enlargement of dark regions. The process of morphological erosion is shown in Figure 2.2.

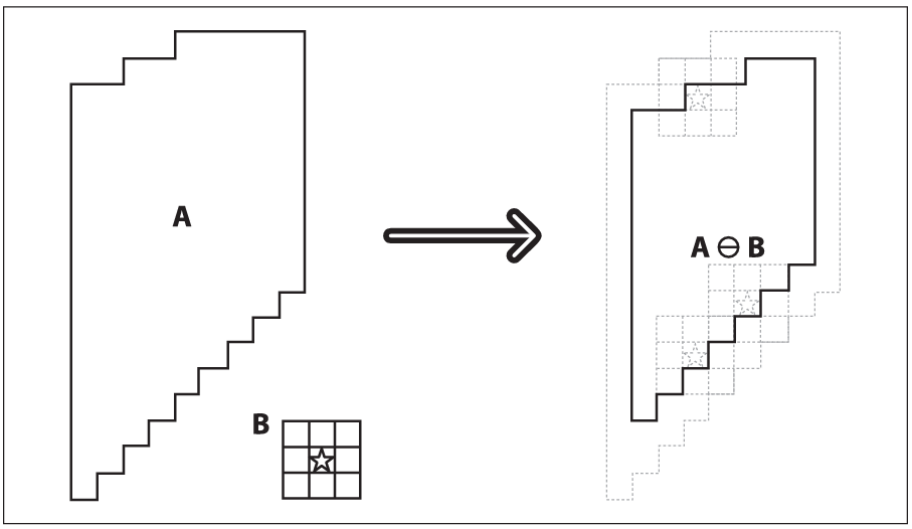

Figure 2.2: Visual representation of the morphological erosion process. Image source: [1].

2.3.4 Morphological Snakes

Morphological snakes build upon Active Contour Models explained in Section 2.3.2, using morphological operations like erosion and dilation, explained in Section 2.3.3. This approach simplifies the curve evolution process by using basic morphological operations to achieve similar results as traditional PDE-based methods.

In morphological snakes, the curve is represented as a binary level set where the interior of the curve is marked with 1 and the exterior with 0. The evolution of the curve is achieved through operations such as dilation and erosion. Dilation expands the boundary of the curve, while erosion contracts it. These operations correspond to the internal force and smoothing operations in traditional active contour models.

The morphological approach avoids the complex numerical implementations and reinitializations required in PDE-based methods. Instead, it uses straightforward inf-sup operators to approximate the differential operators involved in curve evolution. This makes the algorithm faster and more stable.

By applying morphological operators iteratively, the algorithm adjusts the curve to align with significant image features like edges and lines. This method provides an efficient and robust way to adjust the curve for accurate segmentation with reduced computational complexity.

2.4 Evaluation Metrics

This section presents the evaluation metrics used to evaluate the developed segmentation method. In Figure 2.3 below, the union and intersection of two sets, A and B, are displayed. For this project, A & B are different segmentation areas that are compared to each other. All metrics used are in the range 0 to 1 with a perfect value being 1.

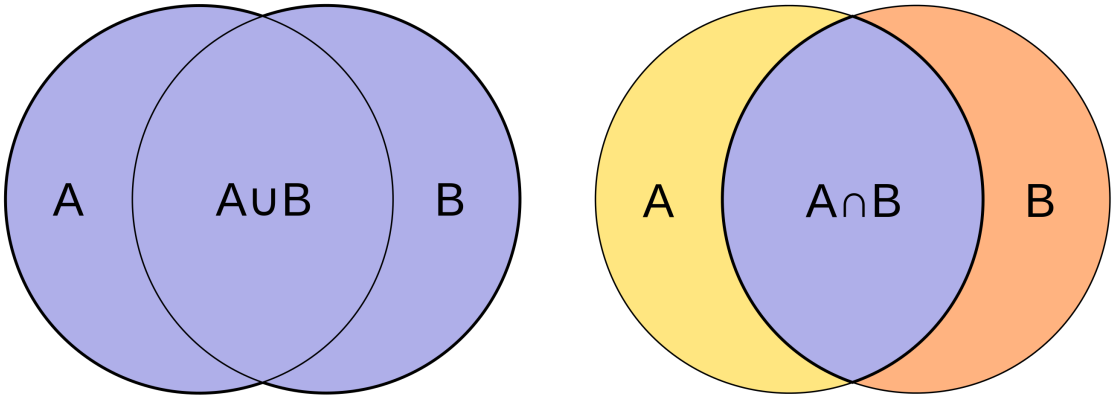

(a) Union of set A and B.

(b) Intersection of set A and B.

Figure 2.3: Images showing union and intesersection of two sets, A and B. Image source: [2]

2.4.1 Dice-Sørensen Coefficient

The Dice-Sørensen Coefficient, also known as the Dice Similarity Coefficient or Dice Score, is a statistical metric used to measure the similarity between two sets [14]. For two sets A and B, it is defined as:

$$DSC(A, B) = \frac{2|A \cap B|}{|A| + |B|}, \tag{2.8}$$

where $A \cap B$ denotes the intersection of A and B (see Figure 2.3b) and $|A|$ and $|B|$ represent the cardinalities (total number of elements) of the sets.

The Dice-Sørensen Coefficient is closely related to the Jaccard index (see Section 2.4.2). The conversion between the Dice-Sørensen Coefficient (D) and the Jaccard index (J) is given by the following equations:

$$J = \frac{D}{2 - D} \quad \text{and} \quad D = \frac{2J}{1 + J}. \tag{2.9}$$

2.4.2 Jaccard Index

The Jaccard Index is another metric used to compare the similarity between two sets A and B, similar to the Dice-Sørensen Coefficient [2]. It is defined as:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}, \tag{2.10}$$

where $A \cap B$ is the intersection and $A \cup B$ is the union of the two sets.

2.4.3 Cardinality Ratio

The Cardinality Ratio is a simple measure that compares the sizes of two sets A and B. It is defined as:

$$CR(A, B) = \frac{|A|}{|B|}. \tag{2.11}$$

3
Methods

In this section, the methods used for this project will be explained.

3.1 Dataset

3.1.1 Description

The dataset utilized in this project comprises several two-dimensional US images stored in JPG format. These images are raw, grayscale scans as displayed on the equipment monitors in the operating rooms. They effectively serve as screenshots of these monitors, also incorporating elements of the graphical user interface (GUI) present during the exam. Each image captures varying extents of the liver tumor intended for segmentation. See Figure 3.1 for reference. In the end, 10 out of these images were chosen based on visibility of the tumor, as well as how the tumors cross-sectional area grew with each slice. This is one of the major limitations with the developed method, which is discussed more in Section 5.2.

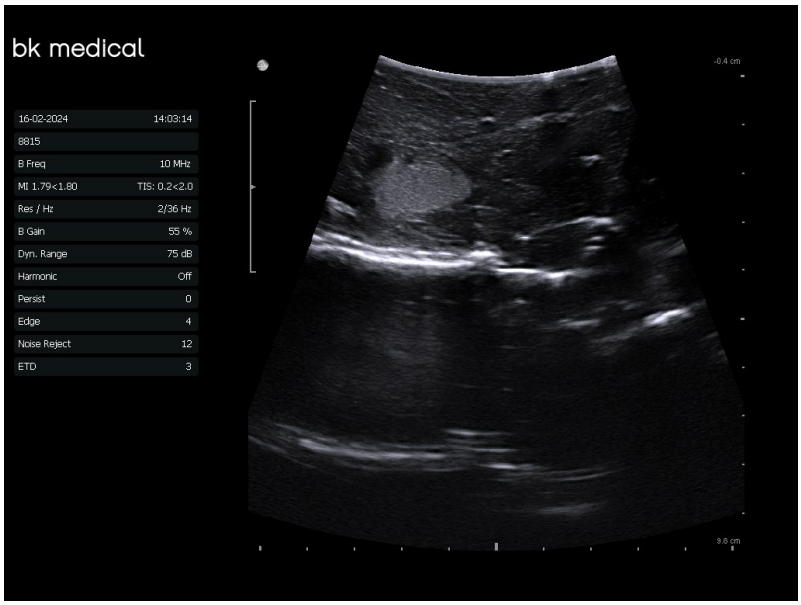

Figure 3.1: Example of an image in the used dataset, where the graphical user interface is seen to the left with the actual US image to the right. The tumor phantom is clearly visible in the top-left part of the image.

3.1.2 Acquisition

The data was collected as part of an experiment conducted by Navari. In this experiment, a tumor phantom was implanted into a bovine liver. During the data collection process, the US probe was tracked in relation to the liver anatomy. This setup allowed for the continuous acquisition of US slices as the probe moved across the liver surface, ensuring a comprehensive representation of the tumor and surrounding liver tissue.

3.2 Preprocessing

The preprocessing of ultrasound images is crucial for enhancing the visibility and distinguishability of the tumor before applying the segmentation algorithm. This section outlines the sequence of preprocessing steps performed, which include cropping, noise removal and closing holes within the region of interest.

3.2.1 Image Cropping

The initial step involves cropping the raw ultrasound images to eliminate any non-relevant parts, such as user interface elements, that could interfere with the analysis. The cropping coordinates were determined manually by identifying pixel indices at two corners of the relevant area. This ensures that subsequent processing operations are exclusively focused on the region containing the ultrasound scan. The result is shown in Figure 3.2.

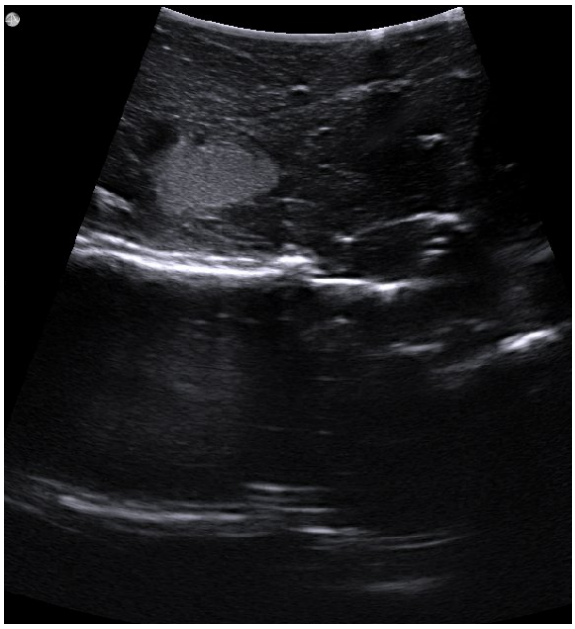

Figure 3.2: Cropped ultrasound image.

3.2.2 Binary Thresholding

After cropping, the next step is to convert the grayscale ultrasound image to a binary format, as illustrated in Figure 3.3. This is achieved using the `cv2.threshold` function from OpenCV, where a threshold value t is specified. Pixels with intensity values above t are set to 1 (white), representing the foreground and those below t are set to 0 (black), representing the background.

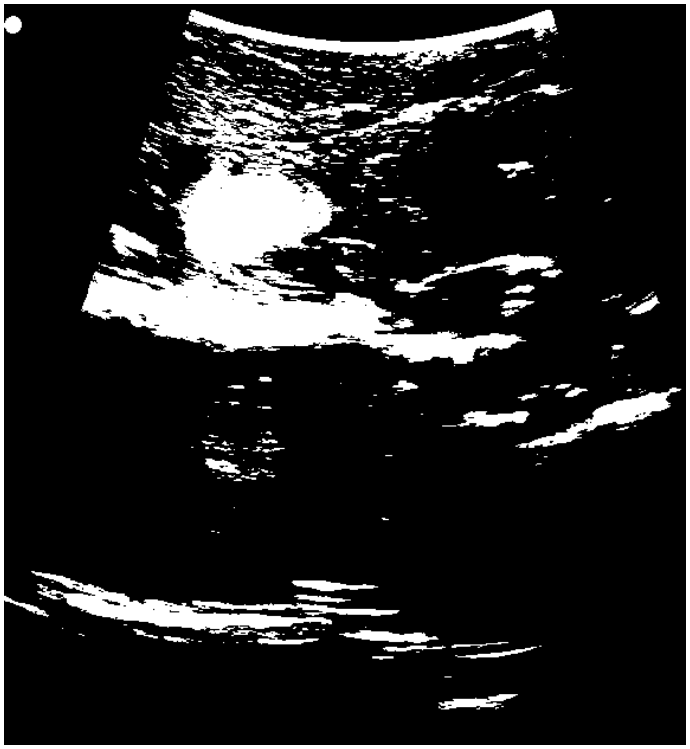

Figure 3.3: Ultrasound image after binary thresholding.

3.2.3 Noise Reduction

Given the presence of significant noise in the binary image, the decision was made to utilize noise reduction techniques. The `cv2.morphologyEx` function is utilized for this purpose, employing morphological operations such as erosion and dilation discussed in Section 2.3.3. The choice of kernel size and the number of iterations were optimized through trial and error to achieve the best results. The outcome of this process is a denoised image where anatomical structures, such as tumors, are more distinct (see Figure 3.4).

Figure 3.4: Ultrasound image after denoising.

3.2.4 Closing Internal Holes

Observations of small internal holes within the tumor structures necessitated an additional preprocessing step. These holes were effectively closed by using the `cv2.morphologyEx` function, reversing the order of operations to apply dilation followed by erosion, unlike the earlier noise reduction step. This technique helps in solidifying the tumor representation, ensuring no internal discontinuities that could affect segmentation accuracy. The processed image, as shown in Figure 3.5, serves as the input for the subsequent segmentation algorithms.

Figure 3.5: Ultrasound image after closing holes.

3.3 Segmentation

This section details the development of the semi-automatic segmentation method, which comprises both manual and automatic components.

3.3.1 Manual Segmentation

The initial phase of manual segmentation involves the surgeon delineating the tumor on the original US image by drawing a polygon. This polygon serves as the initial input for the Morphological Snakes algorithm (refer to Section 2.3.4). Each vertex of the polygon is recorded and used in the preliminary phase of the algorithm. Figure 3.6a displays this step, showing a drawn blue polygon. The accuracy of this initial drawing does not need to be precise, as the Morphological Snakes algorithm is designed to refine this initial estimate effectively.

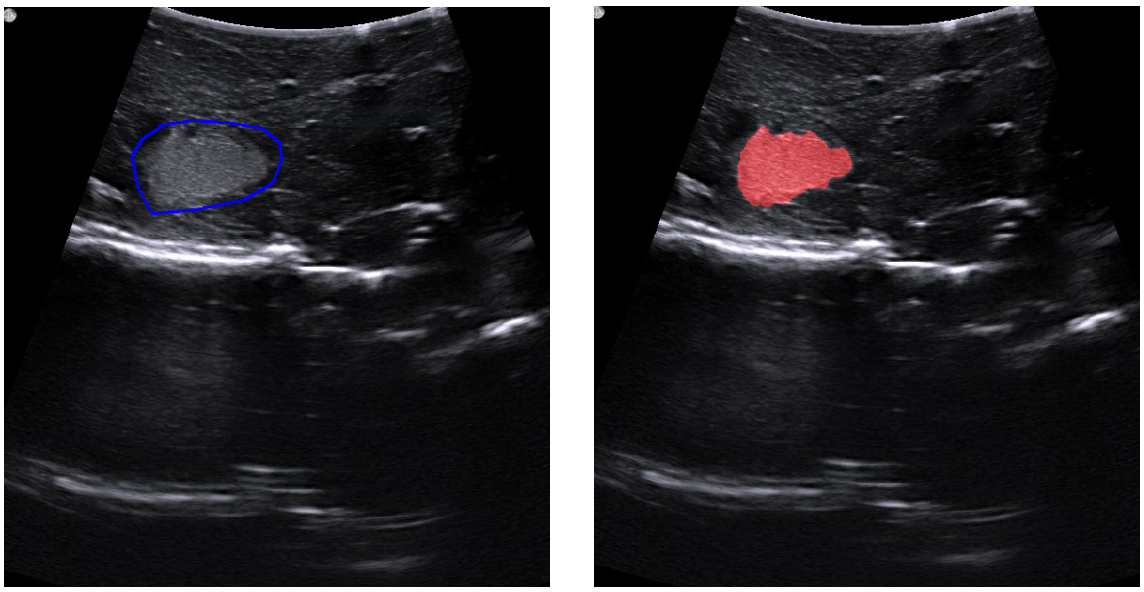

(a) Initial input **(b)** Output segmentation

Figure 3.6: Manual part of the segmentation method with initial guess (left) and output segmentation (right).

3.3.2 Automatic Segmentation

The subsequent segmentation process is entirely automatic. After the initial manual input, the algorithm autonomously segments each subsequent tumor slice. It achieves this by slightly expanding the boundaries of the tumor identified in the previous slice. These expanded boundaries are then used as a baseline for segmenting the current slice, similarly to how the initial manual input is utilized for the first slice. This process is shown in Figure 3.7a.

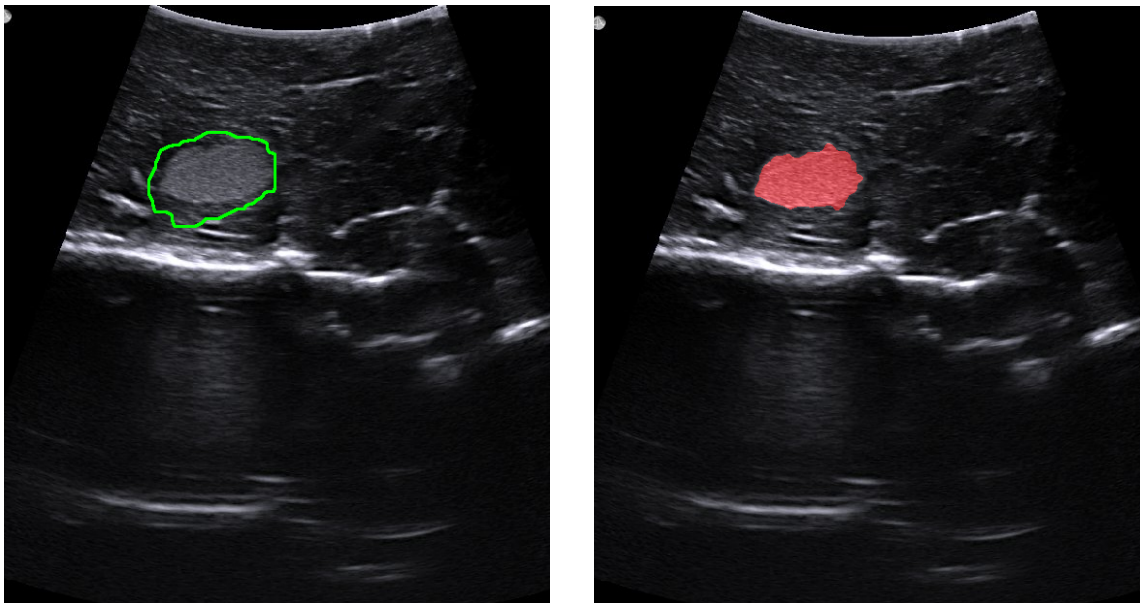

(a) Automatic estimate **(b)** Automatic segmentation

Figure 3.7: Automatic part of the segmentation process with the estimate based on the previous slice (left) and the obtained segmentation (right).

During the segmentation process, the morphological snakes algorithm may incorrectly capture adjacent non-target structures, mistakenly identifying them due to their proximity or similar imaging characteristics. To address this and improve the accuracy of the segmentation, a post-processing step isolates the largest contiguous tumor area. This step incorporates a connected components analysis, typically using a labeling algorithm to distinguish between different regions within the segmented output. Here is how it proceeds:

1. **Labeling of Components**: The algorithm labels each distinct area within the segmented image, differentiating the tumor from other anatomical features or artifacts.
2. **Size Calculation**: It then calculates the size of each labeled area by counting the pixels within each label, identifying the largest region.
3. **Isolation of the Largest Component**: Only the largest component is retained by applying a mask to the segmented image, filtering out all smaller, non-essential areas.

This approach focuses exclusively on the most significant anatomical feature, typically the tumor, enhancing the specificity of the segmentation and streamlining clinical analysis. However, it may overlook smaller or multiple significant tumor sections, a limitation discussed in Section 5.2. This underscores the need for careful interpretation and possible adjustments in cases with multiple lesions or variably sized tumors.

3.4 Desired Workflow Integration

The proposed integration of the developed segmentation method will be thoroughly explained in this section.

3.4.1 Pre-Operative Process

Before the surgery, the patient will undergo CT and MR scans per standard practice. This will enable the surgeon to get an understanding of the shape and location of the liver tumor before the laparoscopic resection surgery. Before the start of the resection, the surgeon will scan the liver surface with the probe, supervising the segmentation process, drawing the initial guess and verifying that the reconstructed 3D model seems to correspond well to their previous understanding of the tumor.

3.4.2 Intra-Operative Process

During the surgery, the surgeon will constantly have the 3D model of the tumor overlaid onto the laparoscopic camera view, enabling accurate visualization and guidance during the procedure. It is also possible to perform new scans at suitable times in order to get an updated 3D model. This could be important due to the deformation that can occur during the surgery.

3.5 Evaluation

It is important to evaluate the developed segmentation method in a suitable way. It was decided that the best way to do this was to compare the different 2D slices obtained by the developed semi-automatic model with the ones from a careful manual segmentation by the author. The statistical metrics described in Section 2.4 (Dice-Sørensen Coefficient, Jaccard Index and Cardinality Ratio) are all used to compare these two segmentations in order to determine the accuracy of the developed semi-automatic segmentation method.

4

Results

This chapter presents the results of the project, focusing on the evaluation of the developed segmentation method and especially challenging slices.

4.1 Evaluation of Segmentation Method

The performance of the segmentation method was evaluated using several metrics. Table 4.1 summarizes the results for individual slices.

Table 4.1: Individual Slice Results

Slice	Dice Coefficient	Jaccard Index	Area Ratio
0	0.56	0.39	0.39
1	0.86	0.76	0.77
2	0.90	0.82	0.89
3	0.92	0.85	1.01
4	0.93	0.87	0.96
5	0.88	0.79	1.09
6	0.93	0.87	1.01
7	0.88	0.78	1.16
8	0.90	0.82	0.82
9	0.60	0.43	0.43

As shown in Table 4.1, the overall segmentation scores are relatively high. However, some slices, specifically slice 0 and slice 9, exhibit significantly lower performance, indicating outliers. These outliers are characterized by significantly lower metric values compared to the other slices. The underlying issues contributing to these outliers are presented in Section 4.2. Slice 6 achieved the best segmentation results, as illustrated in Figure 4.1, where the semi-automatic and manual segmentations show considerable overlap.

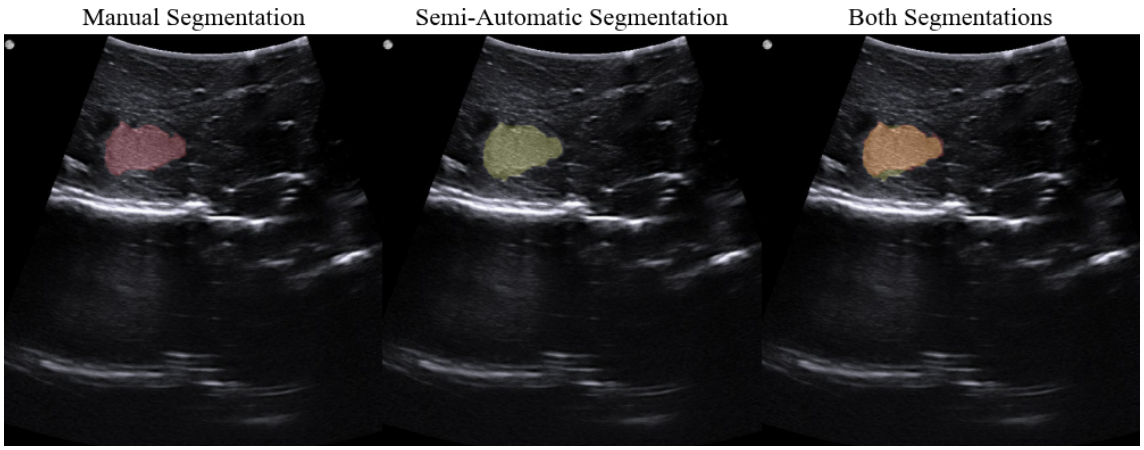

Figure 4.1: Segmentations of slice 6: manual segmentation (left) and semi-automatic segmentation (right).

Table 4.2: Mean Slice Results

Metric	Mean All Slices	Mean Excluding Outliers
Dice Coefficient	0.84	0.90
Jaccard Index	0.74	0.82
Area Ratio	0.85	0.96

The average results for all slices are presented in the first column of Table 4.2. In the second column, the average results are shown if the outliers, slices 0 and 9, are excluded. It is worth noting that the average results are significantly better when excluding these two special cases.

All segmentations can be found in Appendix A.

4.2 Challenging Slices

This section presents the results for the slices that posed significant challenges for the segmentation method.

4.2.1 Multiple Tumor Areas

Slice 9, which contains multiple tumor areas, demonstrates the segmentation method's limitations in handling such cases. As shown in Figure 4.2, the segmentation fails to correctly identify all tumor regions.

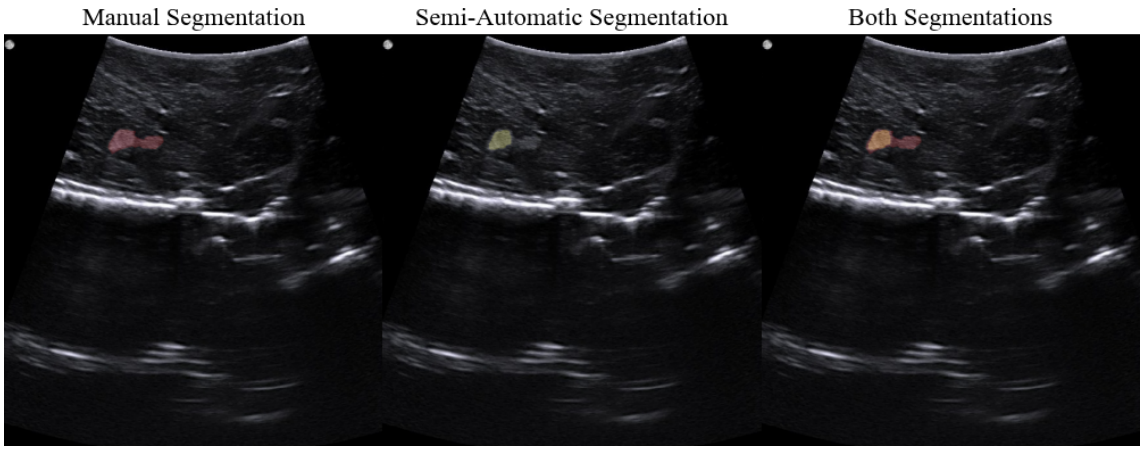

Figure 4.2: Segmentations of slice 9, containing multiple tumor areas, illustrating unsuccessful segmentation.

Figure 4.2 clearly shows that the semi-automatic segmentation method inaccurately segments only one of the tumor areas. To better understand this failure, the binary and final preprocessed images for slice 9 are presented in Figure 4.3. Further analysis of this case is discussed in Section 5.1.1.

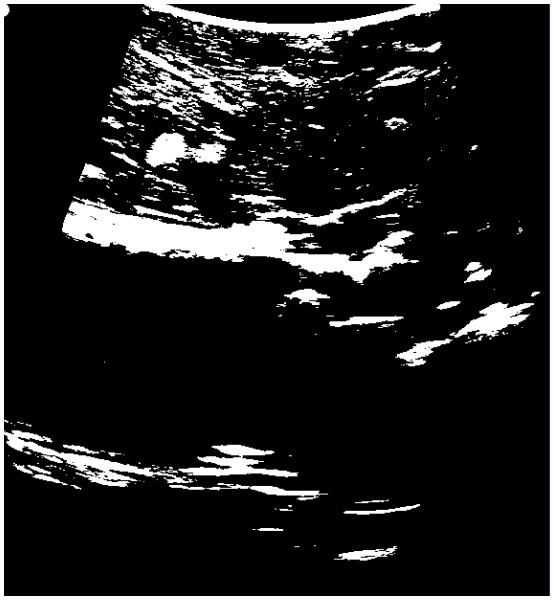

(a) Binary image for slice 9.

(b) Final preprocessed image for slice 9.

Figure 4.3: Binary and final preprocessed image for slice 9.

4.2.2 Insufficient Segmentation Area

Slice 0 shows an insufficient segmentation area, as illustrated in Figure 4.4. The segmentation method struggles with this slice, failing to accurately delineate the tumor region.

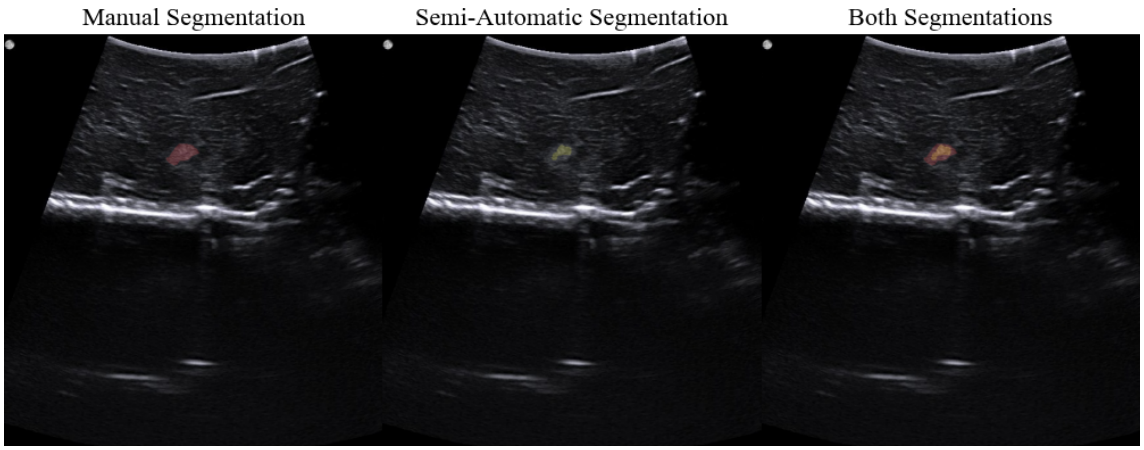

Figure 4.4: Segmentations of slice 0, showing an insufficient segmentation area.

The reasons for this failure can be seen in Figure 4.5, which shows the binary and final preprocessed images for slice 0. The analysis and potential solutions for this issue are further discussed in Section 5.1.2.

(a) Binary image for slice 0.

(b) Final preprocessed image for slice 0.

Figure 4.5: Binary and final preprocessed image for slice 9.

5
Discussion

The results presented in Chapter 4 will be discussed and analyzed in this chapter, along with the project's limitations, possible improvements and future work proposals.

5.1 Challenging Slices

In this section, the results from Section 4.2 are analyzed.

5.1.1 Multiple Tumor Areas

The segmentation method encounters difficulties when dealing with images containing multiple tumor areas. As shown in Figure 4.3b, the tumor appears as two separate areas in the input image for the Morphological Snakes algorithm. A connected components approach (Section 3.3.2) was employed to prevent the segmentation algorithm from mistakenly identifying small non-target structures as part of the tumor, which generally improved the overall results. However, this approach fails in cases with multiple tumor areas, as only one segmentation area is obtained with the current solution. A potential solution to this problem would be to incorporate a toggle that allows the surgeon to specify the number of tumors seen in the segmented slice, thereby improving accuracy.

Another aspect is that thresholding could play a role here. If a more accurate threshold value were chosen, it might result in identifying a single tumor area rather than two. The space between the tumor areas is lower in intensity (see Figure 4.3a), which might be an imaging artifact, as there is only one tumor phantom present in these US scans. Selecting a better threshold value could therefore solve the problem without needing to adjust the number of tumor areas. Further improvements to thresholding are discussed in Section 5.2.

5.1.2 Insufficient Segmentation Area

Analyzing the image in Figure 4.5b, it is evident that a significant portion of the tumor was incorrectly removed, resulting in an undesirably small segmentation area. As seen in Figure 4.5a, the thresholding process works poorly for this particular slice. Significant noise is prevalent and the tumor itself blends in with that noise. This can be improved by adjusting the thresholding step, discussed further in Section 5.2.

This thresholding failure was also combined with an overly aggressive and non-image-specific denoising process. The parameter values, which were determined by trial and error, were uniformly applied to all images. This is a problem, since images might need individual denoising parameters. To address this, other preprocessing methods, tailored to each image could be implemented, something that is discussed more in Section 5.2. Alternatively, providing the end user with adjustable preprocessing parameters (e.g., through a slider) would allow for more precise control and correction when the segmentation area is too large or too small.

5.2 Limitations & Possible Improvements

Several limitations impact the readiness of the developed method. The biggest limitation with the current segmentation method is that it only handles images with high tumor visibility. The US images from the dataset used for this project clearly shows the tumor, which simplifies the segmentation process. For practical application, the method must handle more complex data where the tumor is not as visible, where the gradient between tumor and healthy liver tissue is not as distinct. This would require refining existing preprocessing techniques or developing new ones. One solution could be the use of histogram equalization in order to enhance contrast between tumor edges and surrounding liver tissue before thresholding, something that has not been investigated in this project. Another could be to take advantage of the information from the surgeon's initial guess, beyond just the tumor's location. Automatically cropping the image according to that guess could allow smart techniques to use the pixel information inside (and outside) that area to obtain a good binary image from thresholding. Global thresholding is one option, where the average pixel intensity of the cropped image is calculated and the threshold value is set to that average value. If the image is relatively uniform in intensity, this could be a good option. Otsu's method is another automatic method that minimizes the intra-class variance between foreground and background to find the best threshold value. Adaptive thresholding could be another viable option where different threshold values are used for different regions of the cropped image. All these have not been investigated in this project and could be viable options for improving the segmentation method.

Machine Learning (ML) could potentially be used for both preprocessing the images and the segmentation itself. Firstly, ML could improve thresholding and denoising of the original US images. For instance, if a suitable network were trained on noisy and clean images, it could learn how to efficiently denoise images. This would improve the performance of the developed segmentation method with better input data to the Morphological Snakes algorithm. Secondly, ML could be used for the segmentation part itself as a network can be trained to perform the segmentation. This is a prevalent use of ML with several potential network candidates. One of them is Meta's Segment Anything Model (SAM), a new model designed for segmenting objects that shows promise. This could be a good candidate for segmenting ultrasound images. Another alternative would be to use ML for both preprocessing and segmentation.

The current segmentation method does not allow the surgeon to freely scan the ROI. This is because it is assumed that the US probe is not angled when scanning the object, meaning all US slices are treated as parallel to each other. Allowing an angled probe would provide the surgeon with greater flexibility to scan the liver and tumor, ideally enabling freehand scanning of the ROI. This would save time and enhance the method's utility. It might be possible to remove this assumption and develop a method that allows any angle on the probe. This would likely include the use of projection, where after a slice has been segmented, the edge of that segmentation will be projected in a certain direction onto the next slice to be segmented and then slightly dilated. The challenges here include determining the projection direction and mapping coordinates from one image plane to another. This should be possible and would significantly enhance this method.

Additionally, the order of images is currently manually set in the code, starting with the slice having the largest cross-sectional area and moving outward. This allows the initial guess to propagate smoothly through subsequent slices. The method does not currently accommodate varying slice sizes. In order for the segmentation method to work with more difficult data that can have varying sizes of the cross-sectional area, one would need to be able to monitor and adjust the automatic guesses made during the segmentation process. If a tumor were to unexpectedly grow in the next slice, the guess should be expandable. It is worth noting that in most cases, the tumor is spherical, but ideally, the surgeon should be able to scan the ROI in any order. Using a solution like SAM could be efficient here.

Improvements could also include more input options during and after the segmentation process. The ability to redo the segmentation with a new manual guess for certain slices would be beneficial. Parameter sliders could also be introduced for preprocessing values (discussed in Sections 5.1.1 and 5.1.2). However, it is of utmost importance that the entire workflow remains quick and effortless, with minimal user input. Otherwise, the whole purpose of the solution is undermined.

5.3 Evaluation of Segmentation Method

Selecting a suitable evaluation method is crucial for a project like this. This project compared 2D slices to manual segmentations performed by the author, serving as the ground truth. The reliability of this approach can be debated, as the author is not a sonologist and may not accurately identify tumor edges. Additionally, there is a risk of bias, with the author unconsciously favoring the Morphological Snakes algorithm during manual segmentation. Nonetheless, the tumor edges in the chosen slices are relatively well-defined. Initially, it was planned to have CT scans act as the ground truth when evaluating the developed segmentation method. However, it was eventually deemed unfair and inaccurate to compare two modalities that are so different, which is why manual segmentation was the final choice.

Evaluating the method in 3D, rather than 2D, was considered. This would allow

several other metrics to be utilized, like displacement in X-, Y- and Z-direction, difference in volume, bounding box and center of mass between the two tumor models. It would also be more relevant since the final product utilizes a 3D model. However, due to time constraint this project had to limit its scope to 2D evaluation.

There are also some differences between the metrics used. The Dice-Sørensen Coefficient gives more weight to the presence of common elements between the sets because it counts them twice in the numerator, while the Jaccard Index counts common elements only once, providing a more balanced measure of similarity that considers both shared and unique elements. This makes the Dice-Sørensen Coefficient more sensitive to the presence of common elements and useful when overlap is particularly important, which is extremely important in this project. Given this, it was a good choice to include both metrics, even if the Dice-Sørensen Coefficient might be slightly more relevant. The cardinality ratio, however, is not very useful when it comes to overlap, as it only considers the size of the segmentations. It still serves a purpose as a check to see that the sizes of the segmentations correspond well to each other and could potentially be used as a fail-safe condition, where the system alerts the surgeon if the segmentation area changes significantly compared to the previous slice's segmentation.

5.4 Future Work Recommendations

The discussed limitations and potential improvements present several avenues for future research to enhance the robustness and accuracy of the segmentation method. To build on the findings and suggestions outlined in Section 5.2, future work can be broadly categorized into methodological enhancements, advanced preprocessing techniques and incorporation of machine learning approaches.

Future studies should explore more dynamic and flexible segmentation methodologies, such as implementing features that allow surgeons to specify the number of tumor regions in an image. Additionally, developing an interactive system where users can adjust preprocessing parameters through a user-friendly interface can offer more precise control over the segmentation process, helping to manage variability in image quality and tumor presentation.

Refinement of preprocessing steps is critical for handling more complex data and enhancing segmentation outcomes. Enhanced thresholding and noise reduction techniques tailored to individual images could significantly improve results. Incorporating machine learning-based denoising models can also be effective.

Machine learning (ML) presents significant potential for advancing both preprocessing and segmentation tasks. Training ML models to enhance image quality and implementing advanced ML models like Meta's Segment Anything Model (SAM) can offer substantial improvements in segmentation accuracy and efficiency. Combining ML-based preprocessing and segmentation into a single workflow could also streamline the process and enhance overall performance.

Ensuring the proposed solutions maintain a quick and user-friendly workflow is paramount. While introducing adjustable parameters and manual inputs can enhance precision, it is crucial to balance these additions with the need for a streamlined and efficient workflow. Future research should focus on optimizing user interactions to maintain ease of use. Additionally, developing methods to allow real-time or near-real-time adjustments and feedback can enhance the utility of the segmentation process.

In conclusion, the integration of sophisticated preprocessing techniques, adaptive parameter settings and advanced ML models represents a promising direction for future research. Initial efforts should focus on exploring the use of SAM while retaining key manual preprocessing steps to ensure a balanced approach. By addressing the highlighted areas, future work can significantly enhance the robustness, accuracy and usability of the segmentation method.

For more in-depth information on these future work proposals, see Section 5.2.

5.5 Regulatory Considerations

When developing a solution like this in the highly regulated landscape of medical devices, it is important to keep regulatory aspects in mind. It is non-desirable with a solution that is entirely automatic, which is why this solution aims to be semi-automatic. Regulatory-wise, it is important for the product to have a surgeon being responsible for the decisions made during the surgery. If something were to go wrong, the surgeon should be held accountable and not the product. It would be extremely difficult to release a fully-automatic product onto the market. With this solution, the surgeon makes an initial guess and can monitor the segmentation process, being able to abort if something seems incorrect.

6
Conclusion

The aim of this master's thesis was to develop a semi-automatic segmentation method based on ultrasound imaging. This method is intended to be used as part of the creation of a 3D tumor model for an AR guiding tool during laparoscopic liver tumor resections by Navari Surgical. The project involved utilizing appropriate preprocessing techniques, implementing the segmentation method in Python and evaluating its effectiveness using relevant metrics.

The results obtained indicate that the proposed segmentation method demonstrates promise for its intended application. The initial implementation successfully segmented tumors in ultrasound images, providing a foundation for further development. The method performed well under controlled conditions with images of high tumor visibility, achieving relatively satisfactory results in terms of segmentation accuracy.

However, the method struggled with images containing multiple tumor areas and insufficient segmentation areas due to inadequate thresholding and denoising. These issues highlight the need for more dynamic and flexible segmentation methodologies, as well as refined preprocessing techniques tailored to individual images.

Future work should focus on addressing these limitations through methodological enhancements, such as features allowing surgeons to specify the number of tumor regions and advanced preprocessing techniques to handle more complex data. The incorporation of machine learning approaches, particularly Meta's Segment Anything Model (SAM), presents significant potential. Ensuring a quick and user-friendly workflow while maintaining precision will be crucial for the method's practical application.

In summary, this thesis demonstrates the feasibility of using ultrasound imaging for liver tumor segmentation and provides a solid foundation for further research. The proposed method shows potential for integration into an AR guiding tool for laparoscopic liver tumor resections, contributing to the advancement of surgical guidance technology. Continued research and development, guided by the insights gained from this project, will be essential for realizing the full potential of this segmentation method.

Bibliography

[1] G. Bradski and A. Kaehler, "Learning opencv.," pp. 115–118, 2015.

[2] Wikipedia, *Jaccard Index*. 09 2020.

[3] M. Sharma, P. Somani, C. S. Rameshbabu, T. Sunkara, and P. Rai, "Stepwise evaluation of liver sectors and liver segments by endoscopic ultrasound," *World Journal of Gastrointestinal Endoscopy*, vol. 10, p. 326–339, 11 2018.

[4] U. L. Jayarathne, J. C. Moore, E. , S. E. Pautler, and T. M. Peters, "Real-time 3d ultrasound reconstruction and visualization in the context of laparoscopy," *Medical Image Computing and Computer-Assisted Intervention*, vol. 2017, pp. 602–609, 01 2017.

[5] A. Zygomalas, D. Karavias, D. Koutsouris, I. Maroulis, D. D. Karavias, K. Giokas, and V. Megalooikonomou, "Computer-assisted liver tumor surgery using a novel semiautomatic and a hybrid semiautomatic segmentation algorithm," *Medical Biological Engineering Computing*, vol. 54, pp. 711–721, 08 2015.

[6] I. M. Arias, H. J. Alter, J. L. Boyer, D. E. Cohen, D. A. Shafritz, S. S. Thorgeirsson, and A. W. Wolkoff, *The Liver: Biology and Pathobiology*. John Wiley & Sons, Incorporated, 6th ed., 2020.

[7] S. T. Orcutt and D. A. Anaya, "Liver resection and surgical strategies for management of primary liver cancer," *Cancer Control : Journal of the Moffitt Cancer Center*, vol. 25, 01 2018.

[8] S. Buettner, J. L. van Vugt, J. IJzermans, and B. Groot Koerkamp, "Intrahepatic cholangiocarcinoma: current perspectives," *OncoTargets and Therapy*, vol. Volume 10, pp. 1131–1142, 02 2017.

[9] P. Laugier and Q. Grimal, *Bone Quantitative Ultrasound*. Springer Nature, 2022.

[10] Wikipedia, "Acoustic wave equation," 04 2021.

[11] F. M. Abu-Zidan, A. F. Hefny, and P. Corr, "Clinical ultrasound physics," *Journal of emergencies, trauma, and shock*, vol. 4, pp. 501–3, 2011.

[12] I. Culjak, D. Abram, T. Pribanic, H. Dzapo, and M. Cifrek, "A brief introduction to opencv," in *2012 Proceedings of the 35th International Convention MIPRO*, pp. 1725–1730, 2012.

[13] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *International Journal of Computer Vision*, vol. 1, pp. 321–331, 01 1988.

[14] W. Contributors, "Dice-sørensen coefficient," 05 2024.

A

Appendix

All segmentations for the semi-automatic segmentation method compared to the manual segmentation are displayed below.

A.1 All Segmentations

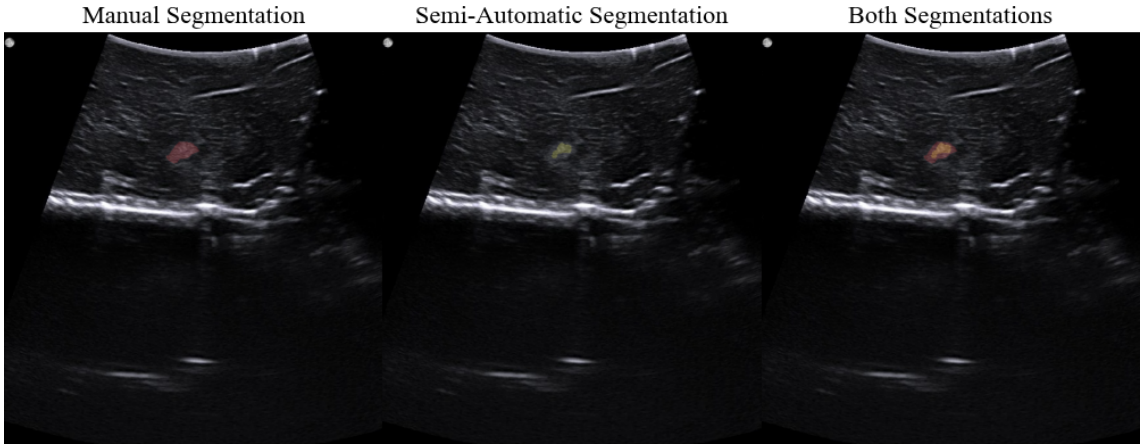

Figure A.1: Segmentations of slice 0, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

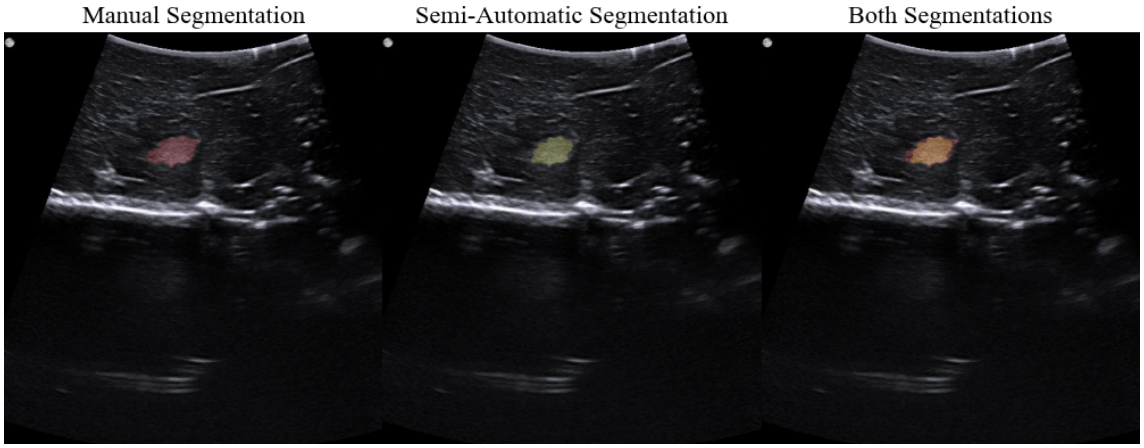

Figure A.2: Segmentations of slice 1, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

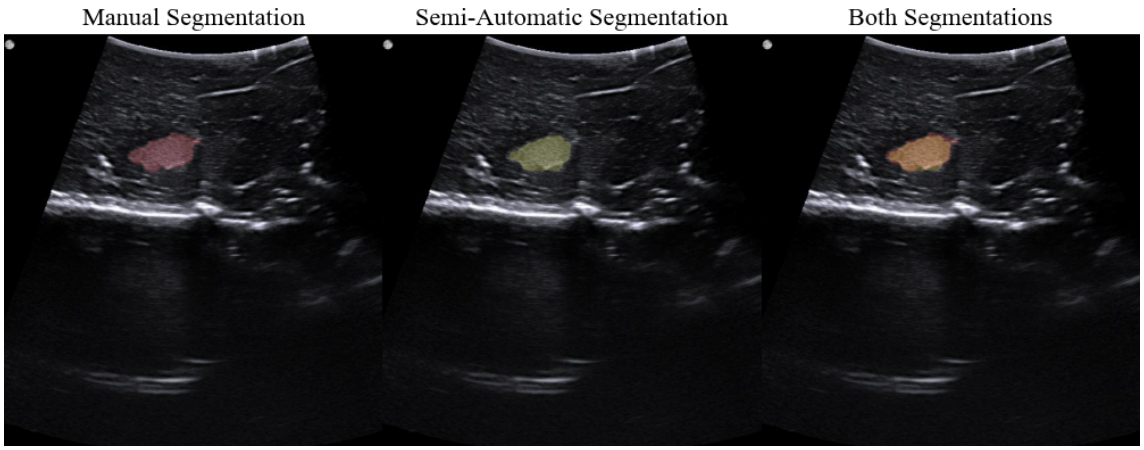

Figure A.3: Segmentations of slice 2, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

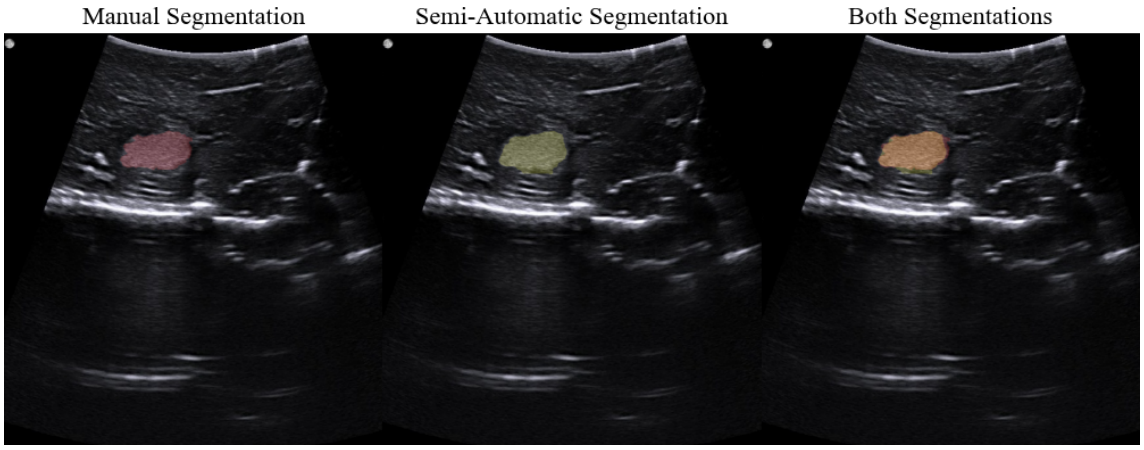

Figure A.4: Segmentations of slice 3, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

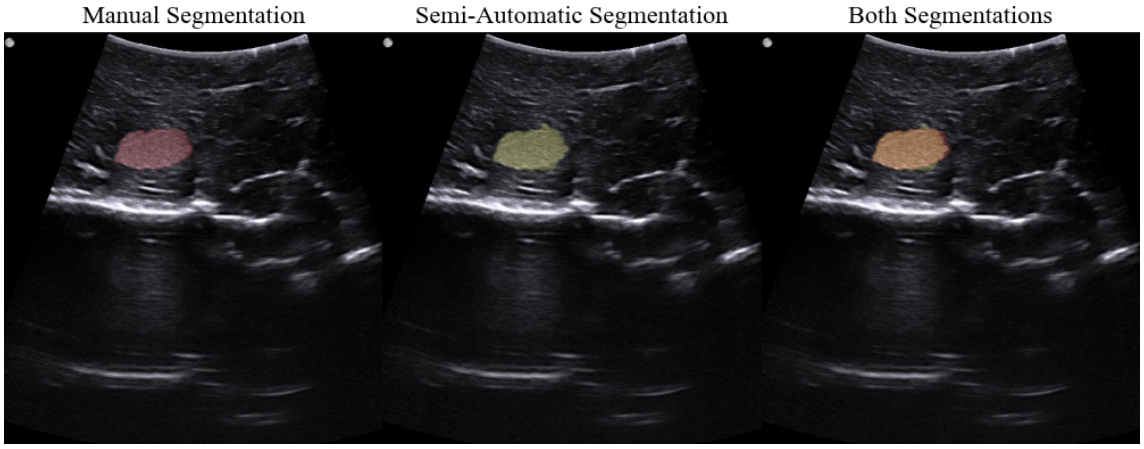

Figure A.5: Segmentations of slice 4, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

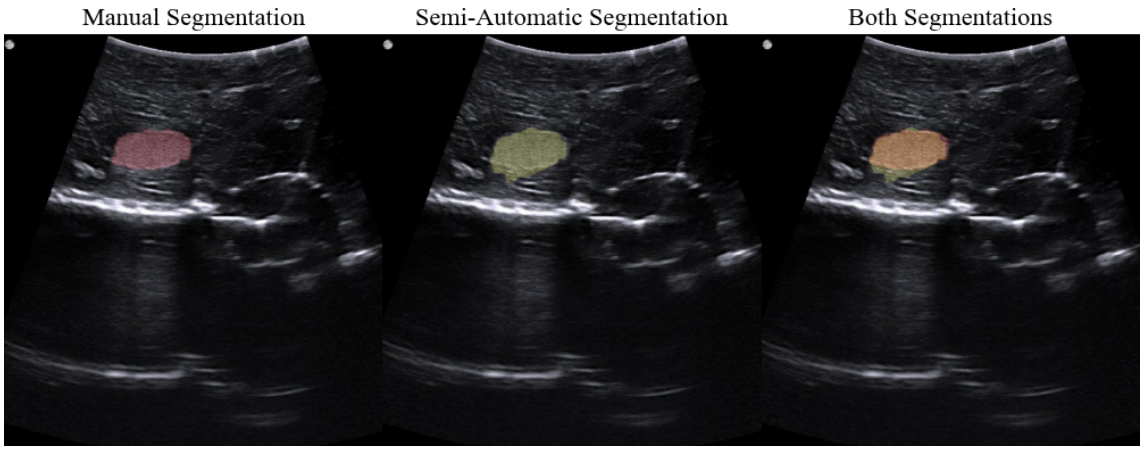

Figure A.6: Segmentations of slice 5, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

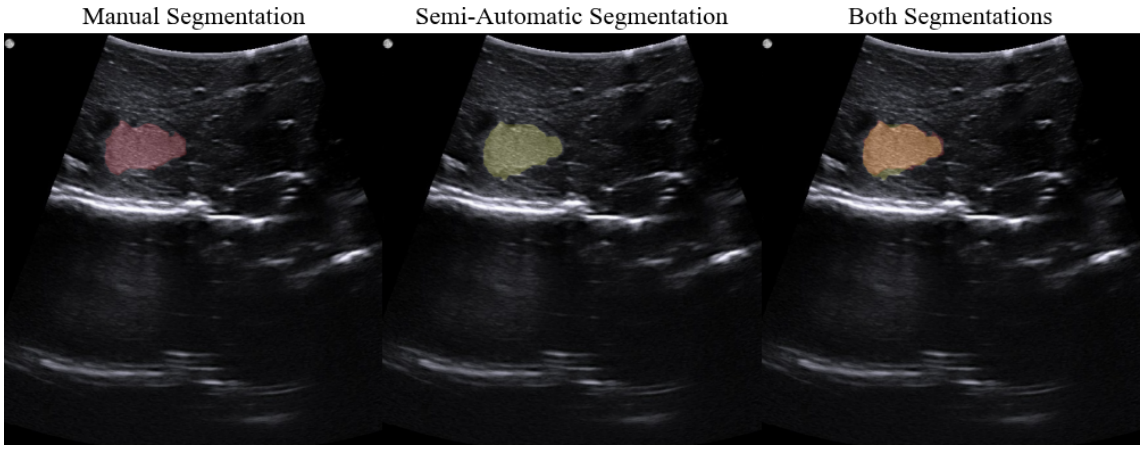

Figure A.7: Segmentations of slice 6, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

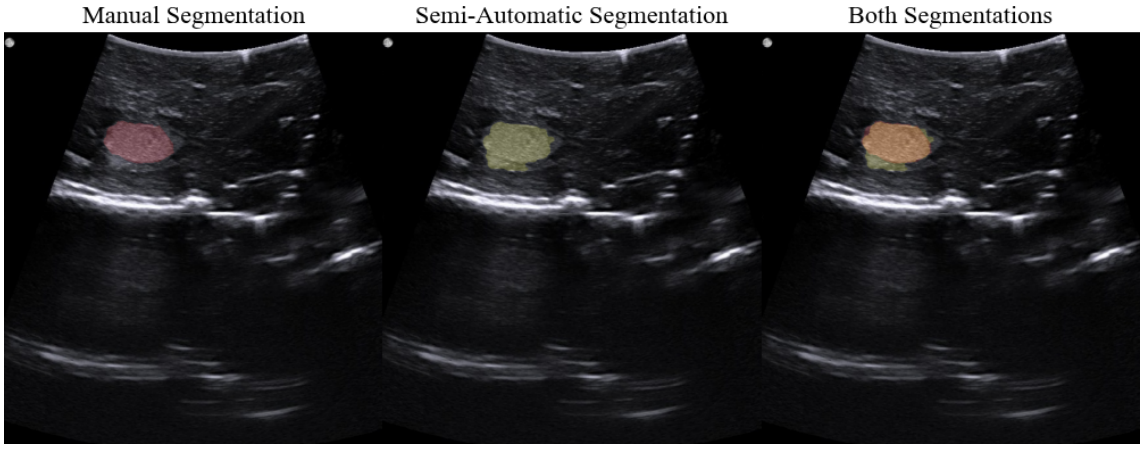

Figure A.8: Segmentations of slice 7, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

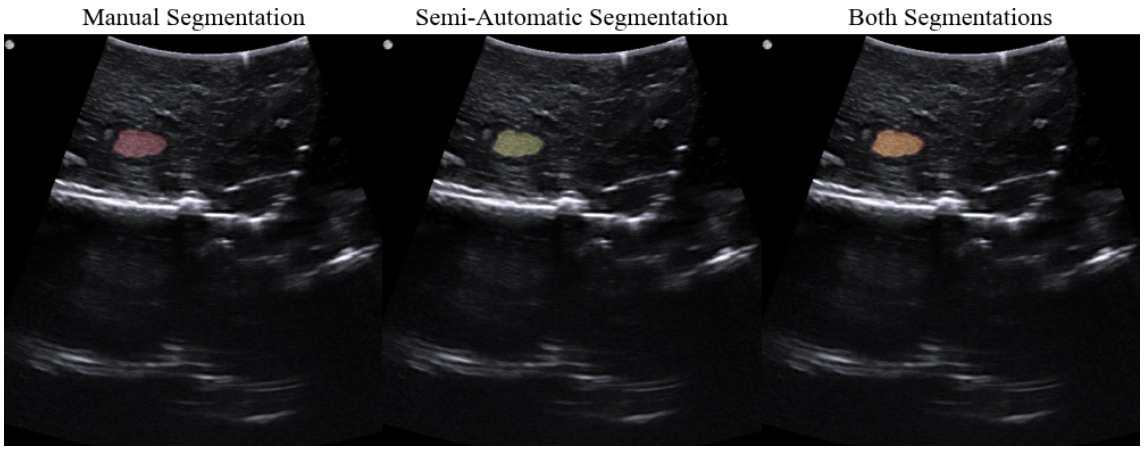

Figure A.9: Segmentations of slice 8, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

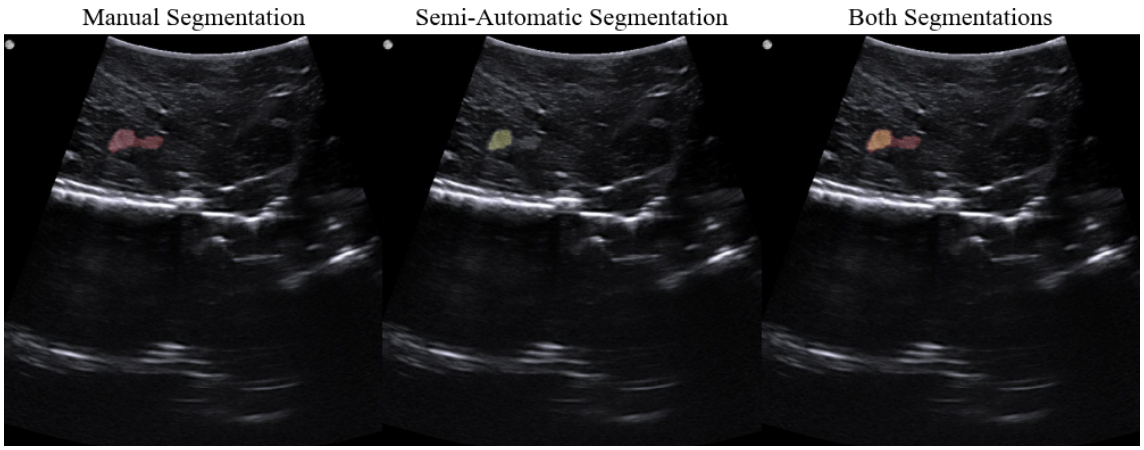

Figure A.10: Segmentations of slice 9, with the manual segmentation method to the left and the semi-automatic segmentation to the right.

DEPARTMENT OF MATHEMATICAL SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY