

RAM

● ROBOTICS
AND
MECHATRONICS

AUTOMATIC SEGMENTATION OF PELVIC FLOOR RING ON 3D ULTRASOUND

E.M. (Enrico) Bertolotti


BSC ASSIGNMENT

Committee:

prof. dr. ir. C.H. Slump
F. Limbeek-van den Noort, MSc
dr. R.M.J. van Damme
dr. H.K. Hemmes

February, 2020

010RaM2020
Robotics and Mechatronics
EEMCS
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands



Contents

1	Introduction	3
1.1	Problem Overview	3
1.2	Ultrasound Data	4
1.3	Problem Approach	4
2	Image Pre-Processing	7
2.1	Module Input	7
2.2	Module Output	7
2.3	Thresholding Filters	8
2.4	Derivative Filters	9
2.5	Morphological Filters	11
2.6	Best Performing Pre-Processing Algorithm	13
3	Position Determination	15
3.1	Module Input	16
3.2	Module Output	16
3.3	Pattern Correlation	17
3.4	Blob Analysis	18
3.5	Best Performing Position Analysis Algorithm	19
4	Angle Determination	21
4.1	Angle Determination By Cross-Correlation	21
4.2	Angle Determination By Edge Curvature Ratio	24
4.3	Best Performing Rotation Analysis Algorithm	26
5	Results	27
5.1	Pre-Processing Results	27
5.2	Position Determination Results	28
5.3	Rotation Determination Results	29
6	Discussion	31
7	Conclusion & Future Work	33
	Appendices	37
A	Global Threshold Filter	39
B	Adaptive Threshold Filter	41
C	Sobel Filter	43
D	Derivative Filter Comparison	45

E	LoG Filter	47
E.1	LoG Filter Standard Deviation σ Variation	48
E.2	LoG Filter after vertical Gaussian blur	49
F	Morphological Operations	51
G	Final Filter Algorithm On Ultrasound Images	53
G.1	Ultrasound Image 1	53
G.2	Ultrasound Image 2	54
G.3	Ultrasound Image 3	54
G.4	Ultrasound Image 4	55
H	Convolution Based Position Determination	57
H.1	Ultrasound Image 1	57
H.2	Ultrasound Image 2	58
H.3	Ultrasound Image 3	58
H.4	Ultrasound Image 4	59
I	Applied Contour Analysis	61
I.1	Ultrasound Image 1	61
I.2	Ultrasound Image 2	62
I.3	Ultrasound Image 3	62
I.4	Ultrasound Image 4	63
J	Correlation Based Angle Analysis	65
J.1	Ultrasound Image 130448	65
J.2	Ultrasound Image 130185	65
J.3	Ultrasound Image 130208	66
J.4	Ultrasound Image 130328	66
K	Slice Based Rotation Analysis	67
L	Pre-Processing Analysis Results	69
M	Position Analysis Results	71
N	Rotation Analysis Results	73

Abstract

A 3-dimensional ultrasound image of a woman's pelvic floor region, in which a pelvic floor ring is present, is analysed for the automatic detection and segmentation of the ring. The segmentation comprises image pre-processing, position analysis and rotation analysis. Several approaches are presented and discussed. The overall best performing result is achieved using a filter chain comprising a 1-dimensional Gaussian blur, LoG filter and a morphological opening operation to pre-process the images, blob detection and analysis to determine the position, and a cross-correlation based rotational analysis. While the location of the ring crosscut can be accurately determined in approximately 70% of cases, the orientation can not be determined with sufficient confidence.

Chapter 1

Introduction

1.1 Problem Overview

The pelvic floor muscles are fundamental in stabilizing the structure of organs and tissue in the pelvic region. However, pregnancy as well as delivery can lead to considerable damage to the muscles, leaving organs in that region (such as the bladder, uterus or rectum) to possibly shift and descend. This is called pelvic floor prolapse. To prevent the serious health problems that can arise from a pelvic prolapse, a pessary (see figure 1.1) ring can be inserted in the vagina to stabilize the vaginal wall[1, 2].



Figure 1.1: Example of a pessary ring

However treatment by pessary has a low success rate of around 60%[1], some patients have complaints and experience pain throughout the procedure. In some cases the ring falls out of the vagina. As not much is known about the position or orientation of the ring, determining the success of a pessary procedure is difficult. 3-dimensional ultrasound data of the patients pelvic floor region was gathered as part of the Gynius¹ project[3]. The goal is to obtain more information about the successes and failures of using a pessary. To do this, the pessary ring is located in the vagina. A first step to locating the ring in the vagina is to locate the ring within the 3-dimensional ultrasound data. Once this has been achieved, the corresponding position of the ring in the human body is established. This work addresses this first step only. In particular, to determine the position of the ring in the data, different algorithms are analysed, presented and evaluated based on their viability for this specific application.

Some particular challenges faced in this project include the high noise level inherent to ultrasound data per se[4], the inconsistent appearance of the ring in the data as well as the partial image of the ring in the given data due the ring diameter exceeding the maximum depth reachable by the ultrasound capture device. Before presenting the general approach taken it is useful to have an overview of the type of data used for the project, more specifically the ultrasound image format.

¹Gynaecological imaging using 3-dimensional ultrasound

1.2 Ultrasound Data

The ultrasound images treated in the Gynius project are generated by a sonic pulse from a 2-dimensional phased array of sound-emitters. At boundaries of different material densities the sonic pulse partially reflects[5]. The intensity of this reflection is based on the difference of material densities - which creates an impedance mismatch - at the boundary and the orientation of the boundary with respect to the source[6]. The distance to the boundary is captured as the time of flight of the sonic signal[6]. Using both the time of flight information as well as the received signal intensity the ultrasound image is reconstructed[7]. This reconstruction is handled by the ultrasound equipment itself and the 3-dimensional data-field, herein also referred to as a 3-dimensional image, forms the input data to this project. As an aid to understand this format, a visualization was created and can be seen in figure 1.2:

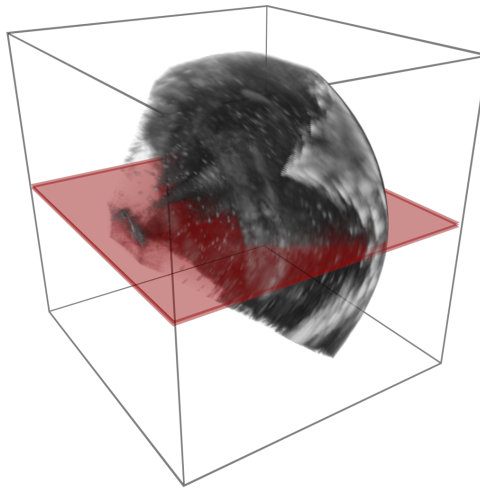


Figure 1.2: 3-dimensional ultrasound image visualization. The red layer is an example of a slice of the data.

1.3 Problem Approach

The general approach taken to tackle this problem comprises the subdivision of the problem into smaller, individual sub-problems. These sub-problems were defined to be the following:

1. Image Pre-Processing
2. Position Analysis
3. Rotation Analysis

We require the solution to each individual sub-problem to fulfill certain, pre-determined targets. Various solutions to the individual sub-problems will be introduced, analysed and their suitability evaluated based on how well the targets are achieved.

Each sub-problem is handled by a module. The image pre-processing module handles the import, cleaning and filtering of the data before it is handed off to be analysed by the position analysis module wherein the coordinates of a recognizable section of the ring are found. Thereafter, the rotation analysis module uses this positional coordinates as a reference point to estimate the rotation of the ring. A visual overview of this process is shown in figure 1.3:

The signal – here the image data – flows from left to right. Each module has an input and an output signal. These signals are labelled in image 1.3 as circles ①, ②, ③ and ④. The module targets are defined for signals 2 - 4, which correspond to the outputs of the modules.




Figure 1.3: Signal Path

This complete segmentation algorithm was developed and optimized on a limited development set of 4 ultrasound images. Afterwards, the algorithm was evaluated on a larger set of 34 3-dimensional ultrasound images. All images used are part of the Gynius project.

Chapter 2

Image Pre-Processing

The aim of the image pre-processing module is not to add information to the underlying image, but to reduce and even remove unwanted noise and artefacts while enhancing the desired features of the image in preparation for the following image analysis algorithms.

2.1 Module Input

The input to this pre-processing stage is a raw 3-dimensional DICOM¹ file and corresponds to signal ① in figure 1.3. An example of a 2-dimensional slice from the input signal is shown in figure 2.1. Prior to processing the data, the data is normalized to the 8-bit range of 0 to 255.




Figure 2.1: Unprocessed ultrasound-image slice

2.2 Module Output

The output of the image pre-processing stage is shown in figure 1.3 as signal ②. To enable a good estimate of the position of the ring, the targets for signal ② are:

1. High contrast of ring edge on background

¹DICOM is an international standard to package medical images[8]

2. Boolean image²
3. Noise free

Noise is here defined as signal intensity variations with small spacial scale and does not include body structures e.g. bones, tissues or veins. The contrast target and the noise level target of the resulting image were evaluated subjectively, rather than quantitatively, due to the time constraints of the project.

The processing consists of a chain of one or more filters. In the following sections, different possible filters are introduced, applied, and analysed. For simplicity only a 2-dimensional slice of the 3-dimensional image will be displayed, however all processing methods are applied slice-wise to the entire 3-dimensional image.

The most effective filter combination is found to be a 1-dimensional Gaussian blur filter followed by a Laplacian of Gaussian filter and finally a thresholding filter.

2.3 Thresholding Filters

While there is a lot of noise in the input image, the edges of the ring are typically visible due to their high brightness in comparison to their surrounding area (see figure 2.1):




Figure 2.2: Unfiltered ultrasound data slice with labelled ring edges.

Therefore a first possible approach is to apply a threshold filter to the image to increase the contrast between the ring edges and the surrounding area. The simplest of the thresholding filters is a global thresholding filter[9].

Global Threshold Filter

A global threshold filter compares the value of each pixel of the image to a set threshold value[10]. If the value of the pixel exceeds the threshold value, the value of the pixel is modified. The thresholding filter chosen modifies the pixels in the following way:

1. All pixel intensities below the threshold value are set to 0
2. The pixels above the threshold value are kept at their original value

²A black and white two-tone image.

This leaves details in the highlighted areas intact, which was thought to assist the image analysis in the later stages. The function of this filter is visualised in figure A.1 (appendix A)

However the results showed that the ultrasound images were too noisy and varied too much to have a threshold filter with a single, global threshold level. Figure 2.3 illustrates the variation of ring edge intensity.

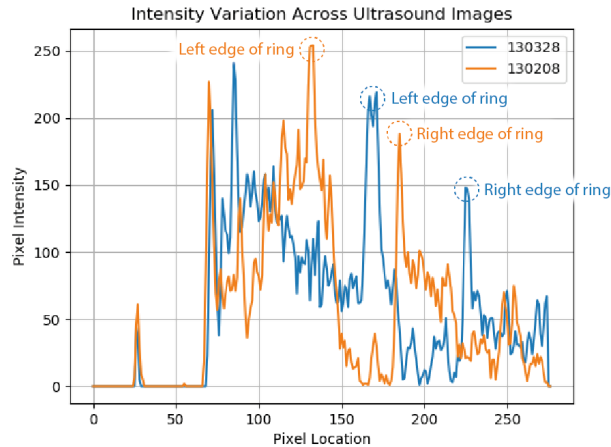


Figure 2.3: Two 1-d slices extracted to illustrate the variation of intensity of the ring in the data. The location of the two ring edges (left and right) of each image ultrasound are labelled

To solve the problem of varying ring edge brightness across different ultrasound images, an adaptive thresholding filter was implemented.

Adaptive Threshold Filter

Instead of using a globally set threshold value (which applies to all pixels equally irrespective of their location), an adaptive thresholding filter compares the current pixel value with the average intensity of its neighbouring pixels[10]. The block-size determines how extensive the neighborhood is around a given pixel. As the filter was applied on 2-dimensional image slices, the block-size is also defined in 2 dimensions. Figure B.1 in appendix B shows the effect of various block sizes of the adaptive thresholding filter on a 2-dimensional slice of ultrasound data.

While the ring edges are more pronounced after applying the adaptive threshold filter, the filter also increases noise and amplifies the neighboring tissue and bone structures. Therefore an adaptive threshold filter is not suitable to reliably filter the ultrasound data.

As simple thresholding filters can not eliminate the noise sufficiently, they fail the noise target and can not deliver the results necessary for the image analysis algorithms. Therefore more complex filters need to be considered.

2.4 Derivative Filters

Compared to thresholding filters - which use the intensity value of the pixel itself - derivative filters use spatial gradient of pixel intensity across the image[11]:

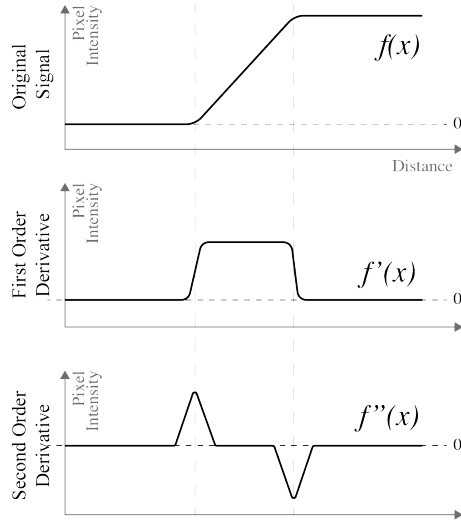


Figure 2.4: Signal representing an edge and its derivatives

This solves the issue of a non constant ring-edge brightness in the ultrasound images. The most commonly used derivative filters can be categorized by the order of the derivative used[11]:

1. First order derivative filters
2. Higher order derivative filters

However all derivative filters amplify noise considerably, increasingly so with derivative order in view of the small spatial scale of the noise. Therefore the image needs to be pre-filtered before the derivative filter is applied. The pre-filter should act as a low-pass filter, preferentially dampening high spatial-frequency noise over large scale structures. The most commonly used low-pass filter is the Gaussian blur filter[11], which is used herein. A 2-dimensional Gaussian kernel function $G(x, y)$ is defined as:

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (2.1)$$

This kernel is convoluted with the image to create the filtered image. Effectively, the combination of the Gaussian filter and the derivative filter acts as a band-pass filter to highlight the medium scale variations.

First Order Derivative Filters

There are many first order derivative filters, the most commonly known being the Sobel and Prewitt filters. These filter differ by the way the gradient is calculated.

make small tables that show each kernel

The results of applying a Sobel filter to the ultrasound image can be seen in figure C.1. The results show that the Sobel filter is not optimal for preparing the ultrasound images for two reasons:

1. At small kernel sizes (less than 5x5 pixels) the contrast of the ring edges with the background is very low
2. At bigger kernel sizes the ring edge disappears in the noise of the image

Both these factors will make it more difficult for the image analysis algorithm to identify the ring position. Other first order derivative filters were implemented (see figure D.1), however they share the same issues, namely sensitivity to noise and lack of contrast around the ring area.

Second Order Derivative Filters

Second order derivative filters respond to the rate of change of the pixel intensity gradient. The simplest 2-dimensional second order derivative filter is the Laplacian operator ∇^2 ,

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \quad (2.2)$$

The Laplacian of Gaussian (LoG filter) combines a Gaussian operator to smooth the image before applying the Laplacian operator[11]. The amount of smoothing prior to the application of the Laplacian operator is determined by the standard deviation parameter. The effect of varying the standard deviation σ of the Gaussian smoothing filter (equation 2.1) on the slice of ultrasound data is shown in figure E.1. There is a trade-off between image artifacts and retaining the shape of the ring edges. The optimal standard deviation was chosen to be $\sigma = 4$ (figure 2.5):




Figure 2.5: Ultrasound image with LoG filter ($\sigma = 4$) applied

The ring edges are clearly visible, however there are a lot of horizontal artefacts in the image. This is solved by smoothing the image vertically prior to applying the LoG filter:




Figure 2.6: LoG filter with vertical Gaussian blur filter signal path

The effects of varying the standard deviation of the vertical Gaussian blur filter can be seen in figure E.2. Similar to determining the optimal standard deviation parameter of the LoG filter, a trade-off exists between retained detail and noise presence in the image. The optimal standard deviation value for the vertical pre-blurring Gaussian blur filter was determined to be $\text{GaussVertical}_\sigma = 7$: Combining the 1-dimensional Gaussian blur filter with the LoG filter, 2 of the 3 targets have been achieved: the result is a boolean image with high contrast between the ring edges and the background. The final step is reduce spatially small noise particles using morphological filters.

2.5 Morphological Filters

Morphological operations are often used to extract image shape features, such as edges, fillets, holes, corners and so on[11]. However, in this case it is going to be used to reduce and eliminate



Figure 2.7: LoG filtered ultrasound image slice pre-filtered with a 1-dimensional Gaussian blur

small spatial variation noise that was not yet filtered out. Since the output of the previous stage is a binary image, binary morphological operators can easily be used.

Two of the most basic binary morphological operations are dilation and erosion. Dilation expands all of the shapes (here shapes are the clusters of white pixels on the black ground) in the image by a certain amount, while erosion shrinks them by a given amount[11]. This process can be seen in figure 2.8:




Figure 2.8: Example of erosion and dilation of a binary image [12]

As the noise in the image is smaller than the outlines of the ring, applying an erosion operation before a dilation operation erodes away the small shapes in the image completely before expanding the remaining shapes back to their original size. Additionally, by dilating them from an eroded state, small crevasses and corners get filled in, creating a cleaner and more smooth shape. The effect of this operation on the ultrasound data is shown in figure F.1.

Choosing the optimal kernel size is dependent on the initial thickness of the ring edges. To ensure that the ring edges do not disappear along with the noise in the morphological erosion process, a kernel size smaller than the thickness of the ring edge (in pixels) needs to be chosen. In theory this implies choosing a kernel size smaller than 6 pixels (for this ultrasound image). However, since the thickness of the ring border is not constant across all ultrasound images, the kernel thickness of the morphological opening was chosen to be 4.

2.6 Best Performing Pre-Processing Algorithm

In summary, a combination of vertical Gaussian blurring, LoG filtering and morphological opening performed the best on the 3D ultrasound data. Figure (Create this figure) shows the output of this filtering approach on other 2-dimensional ultrasound data slices in the patient data set.

Now that a sufficiently well performing filter has been found, a suitable method for obtaining a initial estimate for the position of the ring needs to be found. Up till now the position of the ring has been located by human evaluation. The following chapter presents, implements and analyses different approaches to locate a point on the ring reliably.

Chapter 3

Position Determination

Before the position of the ring can be determined, a feature of the ring needs to be chosen at which the rings position is defined. A feature of the ring which appears consistently across all images is the rings crosscut. The red circle in image 3.1 marks the rings crosscut:

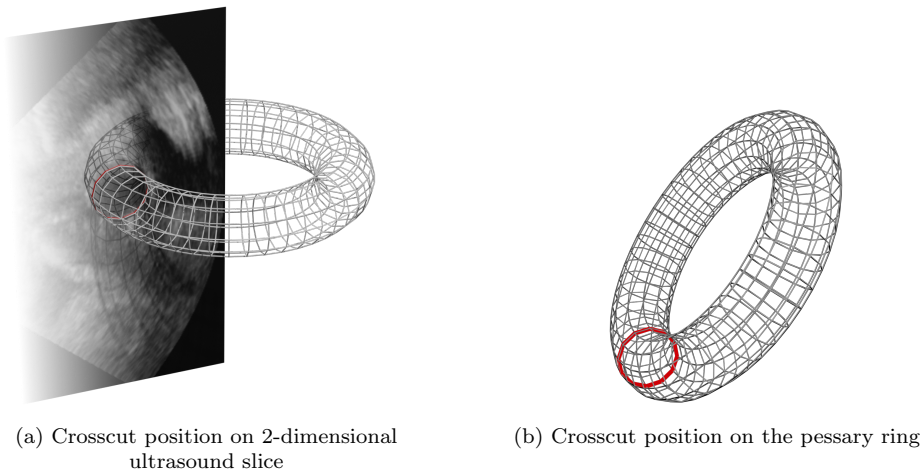


Figure 3.1: Illustration of crosscut position on ring

Figure 3.2 illustrates how the rings crosscut appears in a 2-dimensional slice of the ultrasound image. The position of the rings crosscut is defined by the x-, y-, and z-coordinates of the crosscuts central point.




Figure 3.2: Crosscut appearance and center-point on a 2-dimensional ultrasound slice

In this chapter two different approaches to determine the location of the rings crosscut are implemented and compared.

3.1 Module Input

The input to this module is the raw data (signal ①) in figure 1.3, as well as the filtered image produced by the pre-filtering module.

3.2 Module Output

The output of this module consists of the position of the two ring edges corresponding to the crosscut of the ring. Their locations are each defined by an x -, y - and z -coordinate. An example of the desired output is shown in figure:




Figure 3.3: Crosscut position located at x -position 153, y -position 157 and slice depth 109

3.3 Pattern Correlation

A first approach to determining the location of the ring crosscut is to correlate a image representing the pattern of the ring crosscut with every 2-dimensional slice in the filtered 3-dimensional ultrasound image. The pattern image representing the ring crosscut can be seen in figure 3.4.




Figure 3.4: 2-dimensional correlation pattern

Cross-correlating a 2-dimensional pattern with every 2-dimensional slice of a 3-dimensional array returns a 3-dimensional array of correlation coefficients. The location of the largest correlation coefficient in this array marks the position of the rings crosscut. at which the correlation coefficient is highest marks the position of the ring crosscut in the data. The image used for This was implemented by correlating a 2-dimensional pattern image (see figure 3.4) with every filtered slice in the 3-dimensional ultrasound image. To demonstrate the output of the correlation function, a slice from a filtered 3-dimensional ultrasound image is chosen and shown in figure 3.5. This is the image with which the pattern will be cross-correlated.




Figure 3.5: Filtered slice of ultrasound at z-Depth 104

The values in the cross-correlation matrix result contain both negative and positive values. Negative values indicate a negative correlation, where the inverse of one image is similar to the second image. As negative correlation coefficients have no meaning for the position determination of the ring, they can be clipped to 0. The output of the cross-correlation can be seen in figure 3.6a:

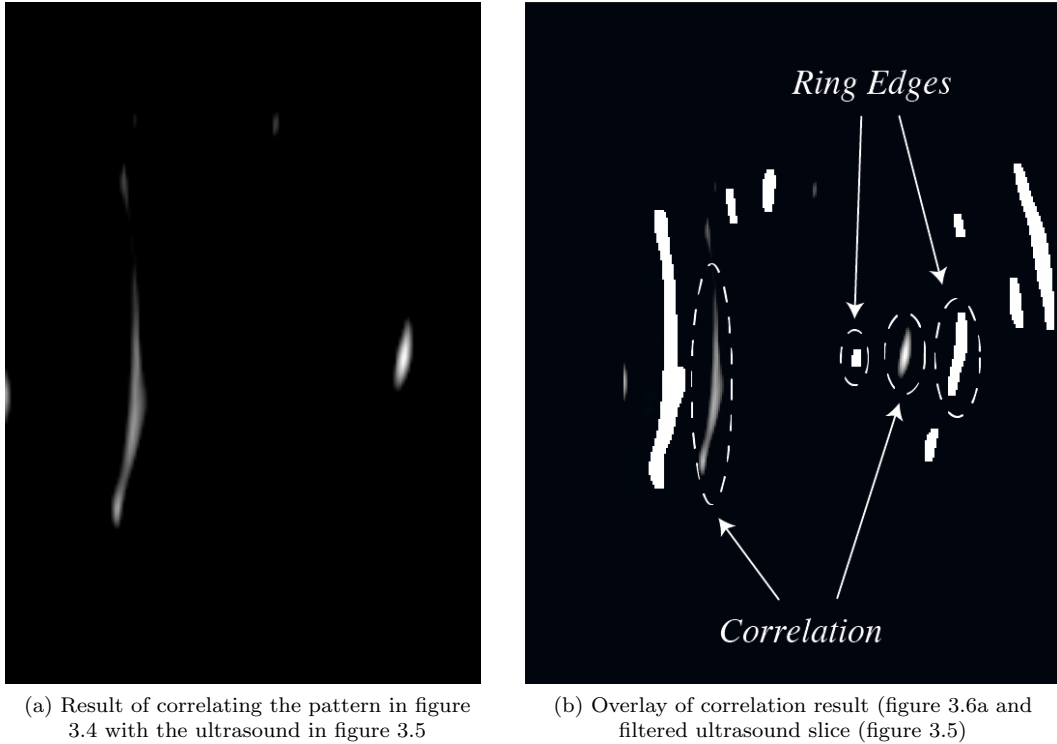


Figure 3.6: Cross-correlation output example

The bright spot on the right shows the point at which the image most matches the pattern. Overlaying the two images shows the bright cluster in the middle of the ring edges (see figure 3.6b).

Applying this technique to the set of development images shown in appendix F, the center point of the ring is correctly located in 3 of them. Figure H.2b shows the false positive that can occur due to the ring in the ultrasound image having a slightly different diameter to that used to create the search pattern.

A solution to this is to darken specific parts of the slices based on the probability of the ring occurring at that position before cross-correlating. For example this could mean darkening the slices based on their distance to the center slice of the 3-dimensional ultrasound data, where based on observations the ring is most likely to be.

Using these assumptions makes the convolution results more stable and less susceptible to false positives, however it also reduces the robustness of the entire solution as any data not conforming to the assumptions cannot be detected.

Therefore a different approach is attempted. Instead of searching for a pattern in the image, all shapes and their relations to one another in the filtered ultrasound image are analysed and their probability of being the ring being calculated.

3.4 Blob Analysis

Each filtered ultrasound image slice (2-dimensional) is composed of two elements: the black background and small clusters of white pixels (hereafter referred to as blobs). Amongst all blobs, there is a blob pair where the two blobs represent the visible edges of the ring crosscut. To identify this blob pair, use is made of the following blob pair properties:

1. The two blobs are 60 pixels¹ apart, which corresponds to the diameter of the ring crosscut. This diameter remains constant across pessaries with varying main ring diameters.
2. A line joining the two blob centers is predominantly horizontal
3. The pair of blobs are assumed to be near the center of the 3d ultrasound image

Similar to a machine learning algorithm, each of the properties listed above is assigned a weight. For each image slice, every possible connection between two blobs on the same slice is generated. The deviation of the connections properties to those listed above is calculated and weighted. This produces a numerical value between 0 (unlikely) and 1 (certain) for each connection. The connection with the highest value corresponds to the modules best guess of the ring crosscut.

Figure 3.7 shows a result of this algorithm. The probability of each connection between two blobs is reflected in its color, with red signifying the connection is very unlikely to be part of the ring, while green signifies a high chance that the connection corresponds to the ring crosscut.




Figure 3.7: Connections between contours are visualized with lines, the colors illustrate the probability of the connection being the ring crosscut (green: likely, red: unlikely)

All 4 of the ring crosscuts were successfully and accurately located in the ultrasound image development set. These results can be seen in appendix I.

3.5 Best Performing Position Analysis Algorithm

Of the two approaches presented, the blob analysis based algorithm is less sensitive to variations in the appearance of the crosscut and gives more information about the shapes of the ring crosscut. Of the 4 images to test the algorithms, the blob based algorithm correctly locates the crosscut in all 4 cases while the correlation based algorithm solely locates 3 of the 4. Therefore the blob analysis algorithm is the preferred algorithm to determine the position of the ring crosscut.

¹The pixel value is a function of the ultrasound capture resolution as well as the rings physical dimensions, both of which stay constant across the Gynius ultrasound data set.

Chapter 4

Angle Determination

Once the location of the ring crosscut has been identified, the orientation of the ring needs to be determined. There are two possible axis along which the ring can rotate. These are depicted in figure 4.1:

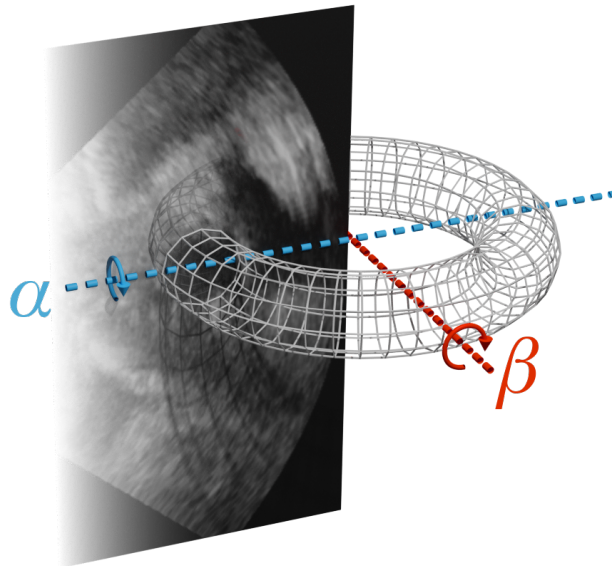


Figure 4.1: Labeled Rotational Axis Visualization

The ring is considered symmetrical along the third axis (perpendicular to both the α - and β -axis). Therefore the rotation around this axis does not need to be determined. Due to the time constraints of the project the rotation analysis was limited to analysing the ring rotation around the α -axis. To determine the rotation along the α -axis, two possible methods are implemented and tested: a correlation based approach and a slice comparison based approach.

4.1 Angle Determination By Cross-Correlation

This correlation based approach is similar to the correlation based method used to determine the position of the ring described in section 3.3. However, instead of varying the position of the pattern, the rotation of the pattern is varied.

The search pattern is correlated image with a 2-dimensional slice of the ultrasound image at the closest point of the ring to the ultrasound probe. This slice is perpendicular to the α -rotational

axis. The position and orientation of this slice is shown in figure 4.2:




Figure 4.2: Orientation of the slice used as a base image for correlation. The slice is shown in green.

Figure 4.3 shows the images corresponding to the green slice shown in figure 4.2 for ultrasound image 130448. The ring shape is difficult to recognize in the unfiltered ultrasound data (figure 4.3a), however after applying a LoG filter to the image the α -rotation of the ring becomes more clear (figure 4.3b).

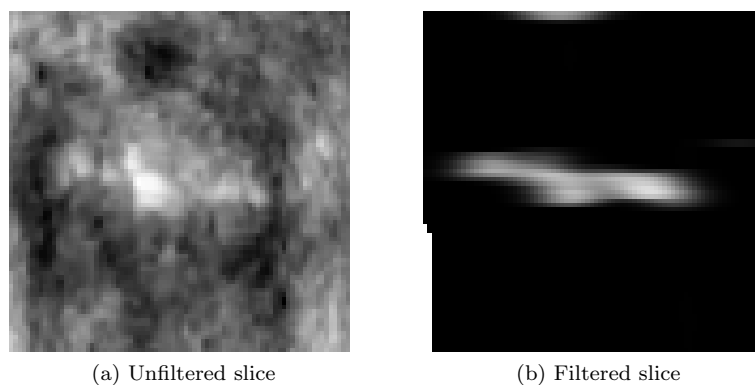


Figure 4.3: Slice showing the top of the ring in the ultrasound data

From this perspective the shape of the top of the ring resembles an ellipse. As this observation is consistent with the other images in the development data set, the pattern that is cross-correlated with the slice is chosen to be an ellipse shape. Examples of this pattern and its rotations is shown in figure 4.4:




Figure 4.4: Examples of patterns that are cross-correlated with the ultrasound slice to determine the rotation.

To cover all possible ring rotations, the pattern is rotated from -90° to 90° in steps of 1° . At each rotation the pattern is cross-correlated with the image and the maximum correlation value stored as a function of rotation. The result of the cross-correlation operation is shown in figure 4.5:

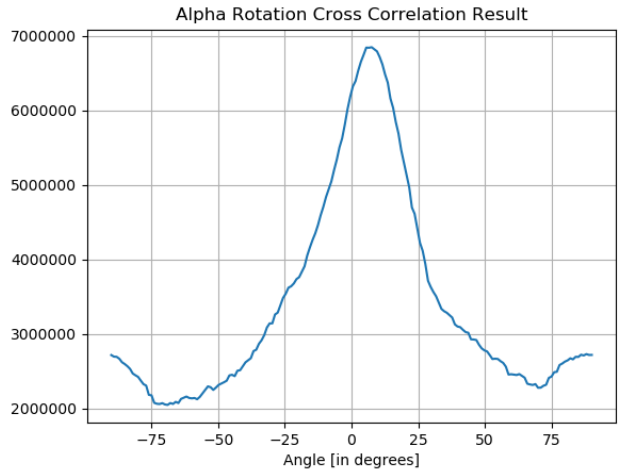


Figure 4.5: Cross-correlation α -rotation analysis plot for ultrasound image 130448. The peak indicates the maximum correlation between the pattern and the filtered ultrasound image.

The plot shows a clear peak at 7 degrees. This means that the ring is rotated by 7 degrees clockwise relative to the horizontal axis. To see if this angle is indeed correct, a line was drawn at that angle onto the original (figure 4.6a) and filtered (figure 4.6b) image:




Figure 4.6: Rotation angle estimate visualized with line

In this case the method proved to be accurate. However, when applied to all images in the development data set, this algorithm has a success rate of 50% (see figures in appendix J). This

cross-correlation based algorithm is very dependent on the ring being clearly visible and bright from the front angle. The second approach uses the curvature of the different radii to approximate the rotation of the pessary ring.

4.2 Angle Determination By Edge Curvature Ratio

This approach uses two perpendicular slices at of the ring edge to approximate the ring rotation. The layout of these two slices can be seen in figure 4.7.




Figure 4.7: Orientation of slices for curvature based rotation determination

These two slices both show a crosscut of a edge radius of the ring. If the ring has a rotation around 0° , the purple slice (from now on called the horizontal slice) in figure 4.7 will show the curvature of the main, large edge radius of the ring. The orange slice (from now on called the vertical slice) will depict the smaller edge radius which represents the crosscut edge radius of the ring. An example of these two crosscuts is given in figure 4.8:




Figure 4.8: Perpendicular slices of an ultrasound image at the top of the ring

The large ring diameter is visible on the horizontal slice (figure 4.8b), while the vertical slice shows

the crosscut of the ring edge (figure 4.8a).

If the ring is rotated 90° along the α -axis, the vertical slice will show the large edge radius of the ring and the horizontal slice will show the ring crosscut. The size of each edge radius on each slice has consequently been swapped. The same principle can also be applied to smaller rotations of the ring between 0° and 90° . In this case, one of the edges on a slice will have gotten smaller by a certain amount, while the other edge on the other slice will have grown by a certain amount. Consequently, by measuring the ratio of one edge radius on a slice to the edge radius on the other slice, the angle of the ring can therefore be estimated.

A method of comparing the radii of the two ring edge crosscuts needs to be found. As the ring edge appears bright, the slice with the bigger ring edge should be brighter than the slice with the smaller ring edge as a bigger edge implies more bright pixels.

The noise and varying appearance of the ring edge need to be considered. To reduce the impact of noise and varying ring appearance, multiple slice pairs are created and analysed for the same image. To create these slice pairs, the ultrasound image is rotated 1 complete revolution around the α -axis in steps of 1° . At every degree a horizontal and vertical slice is extracted.

Once the set of 360 slice pairs is created, in each slice pair the brightness of the horizontal crosscut is divided by the brightness of the vertical crosscut. A peak in the resulting data will mean that the brightness ratio of the horizontal slice to the vertical slice is maximal, indicating that the ring is perfectly level in the data. As the horizontal and vertical crosscut will reach their peak brightness twice during a full rotation, the data will have two peaks and two valleys. Figure 4.9 shows the result of this theory applied to an ultrasound image:

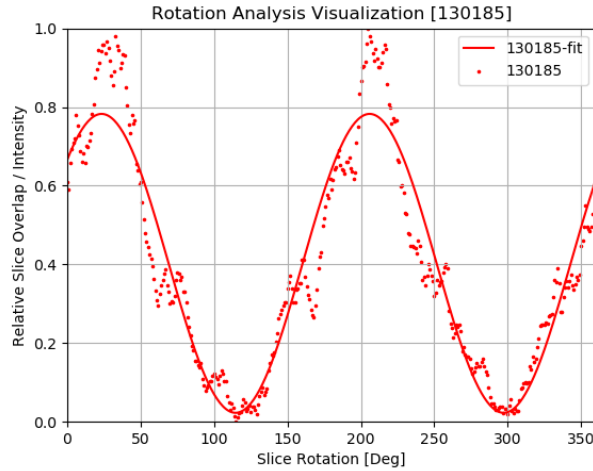


Figure 4.9: Normalized plot of brightness ratio as a function of slice rotation with fitted sin-function.

To reduce the noise-based variations in the data, a sine function of 2 periods in length was fit to the data using the least squares method. The phase offset of the fitted sine function is then used to estimate the ring rotation. The relation between the phase offset and the ring rotation is given by the function:

$$\text{sin}e_{peak} = \text{sin}e_{phase} - 90^\circ \quad (4.1)$$

in which $\text{sin}e_{peak}$ is the location of the maxima of the sine function and $\text{sin}e_{phase}$ is the phase offset of the sine function in degrees. The result of this analysis on the ultrasound image 130185 (illustrated in figure 4.9), shows the first peak at 7 degrees.

4.3 Best Performing Rotation Analysis Algorithm

Both algorithms presented to determine the rotation around the α -axis can not reliably determine the rotation accurately, due to the high noise level and low resolution of the slices. Comparatively the convolution based approach is a more reliable and accurate algorithm and is therefore chosen as the preferred algorithm.

Chapter 5

Results

The best performing segmentation algorithm, as determined from the above work, comprises the following solutions to each module:

- Pre-Processing:** Chain of a 1-dimensional Gaussian blur filter, LoG filter and a morphological opening operation
- Position Determination:** Blob analysis algorithm
- Rotation Determination:** Cross-correlation with rotating pattern

In the following, the results of each module on the 34 ultrasound image are analysed and evaluated.

5.1 Pre-Processing Results

The pre-processing module is considered successful if the ring edges are visible on both the input image as well as the output image. If the edges are visible on the original image, but not on the resulting image, it is considered a failure. If the edges are not visible on the original image, the success is unsure for that image. A different segmentation algorithm should be used for these images. The results of the pre-processing module on individual images in the full 34 image data-set can be seen in figure L.1 in appendix L. Figure 5.1 shows a bar graph representation of the results:

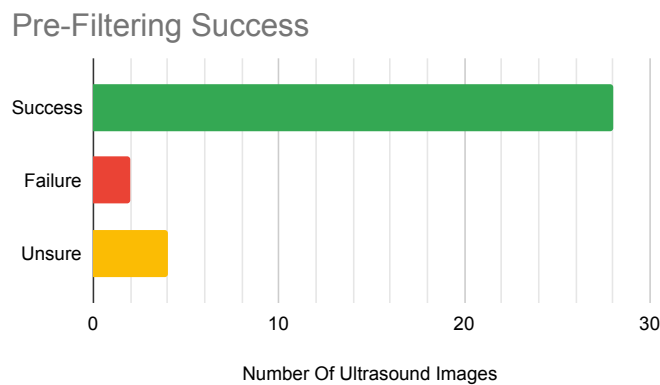


Figure 5.1: Bar-graph representation of pre-filter success rate

The unsure cases are not considered in the evaluation of the success of the module. The pre-processing algorithm successfully filters 28 out of the 30 cases in which the ring crosscut is visible on the unfiltered image. However in the remaining 2 cases, where the ring edge is not sufficiently visible in the data, the ring edge is filtered out. These two cases are shown in figure 5.2a and figure 5.2b:

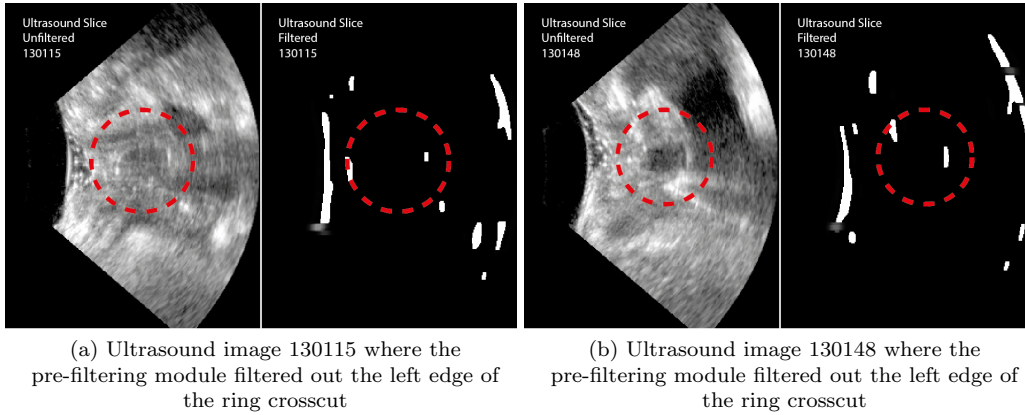


Figure 5.2: Pre-filtering module failures

The left edge disappears due to the pre-blurring using both the 1-dimensional vertical Gaussian blur and the 2-dimensional Gaussian blur. However reducing the intensity of the pre-blurring (by lowering the standard deviation σ of these functions re-introduces noise into the filtered image.

5.2 Position Determination Results

For the position determination module, the definition of a success and failure of the position determination module are given in table 5.1:

Category	Definition
Success	The correct ring edge contours are identified and located
Failure	A different edge pair is identified to be the ring.

Table 5.1: Categories of results from the position determination module

The 6 of the original 34 cases in which the pre-filtering module did not successfully filter the image - therefore automatically resulting in a failure of the position determination module - are not considered as they are not failures of the position determination algorithm per se. The detailed table of results is shown in figure M.1. Figure 5.3 summarizes the results of the position determination module in a bar graph:




Figure 5.3: Bar-graph representation of position determination results

In 19 of the remaining 28 cases the ring position was accurately determined, giving a success rate of approximately 68%. The following case - ultrasound image 130232 - is an example of the position determination module failing to correctly identify the ring. Figure 5.4 depicts both the correct slice depth at which the ring crosscut is visible as well as the wrongly identified ring crosscut.




Figure 5.4: Misidentified ring position in an ultrasound image

This example shows that the properties chosen for the connections are not ideal. The choice of a horizontal connection property is only applicable if the ring is centered vertically in the image. The connection should instead be weighted based on its how parallel it is with the sound emission direction of the ultrasound probe.

5.3 Rotation Determination Results

If the position determination was successful, the rotation of the ultrasound image was analysed. To determine if the accuracy of the rotational analysis, the rotation of the ring in each ultrasound image was manually determined. The difference between the manually determined angle and the angle determined by the both the cross-correlation based approach as well as the slice based approach can be seen in the table in figure N. Figures 5.5a and 5.5b summarise the errors:




Figure 5.5: Error of rotation analysis modules

The results are categorised based on their error and counted. Figure 5.5a shows that using the correlation based analysis, many of the results have an error less than 11° . However figure 5.5b shows that only 5 results obtained using the slice analysis algorithm had an error of less than shows that results had an error of less than 12° .

Neither the cross-correlation based approach nor the slice based approach can deliver a consistently accurate estimate for the α -rotation of the ring.

Chapter 6

Discussion

The results show a partial success of the algorithm. A likely cause of this partial success is the limited number of images in the development set. The lack of variation associated with this small set of test images possibly led to a biased optimisation of the algorithms parameters. This bias can probably be reduced or eliminated if the parameters are optimized over a larger set of images. Unfortunately, the time constraints imposed on this project did not allow for a more comprehensive image set for development. Nevertheless, useful insights have been obtained.

In approximately 68% of the 28 cases, the position determination module successfully identified and outlined the ring crosscut. In the remaining 32% of cases, a structure not part of the ring was identified as the ring crosscut. This shows that the weighting of the parameters used to weight the connections are not optimal. In addition to the euclidean length of the connection and the angle relative to the horizontal axis, there are many additional properties of the contours and connections that could be weighted to more robustly identify the crosscut of the ring. These can include, but are not limited to: the relative size of each contour at the ends of the connection, the x-, y-, and z-coordinate of the connection, as well as the similarity to contours and connections in adjacent slices of data. However due to the sheer number of weighting possibilities of these factors, it would be impractical to determine these manually.

The inaccuracy of the rotation determination algorithms is most likely due to the high amount of noise present in the ultrasound image. A possible solution is to threshold the slices prior to evaluating the brightness. However, since thresholding an image is a non-linear function, this leads to non-linear shifts in brightness. This can negatively impact the quality of the fitted function.

To accurately determine the ring orientation, a different approach, such as a traditional machine learning (ML) algorithm, can be explored to accurately determine the rotation of the ring. However, traditional ML algorithms require a large amount of training data before producing an accurate and reliable result. The labelling of this big data set would require a significant time investment to implement. However, this labor intensive process can be made more efficient by using the position of the ring as a first estimate, and only correcting the data if needed. Therefore, a solution integrating both this initial estimate and a machine learning algorithm should be explored.

When viewed as a whole, the segmentation algorithm can automatically segment the ring in approximately 30% of the cases. A successful segmentation is here defined as identifying the ring crosscut correctly and determining the ring rotation with an error less than 8° . This low general success rate can be largely attributed to the propagation of errors throughout the segmentation algorithm. If the pre-processing module fails, all following modules have little or no chance of successfully segmenting the ring. The effect of error propagation can be diminished by employing multiple parallel approaches, each calculating a confidence metric for its result. By taking the individual confidence metrics into account, wrong results and their error can be filtered out. Due to the low complete success rate, a solution based solely on the approach taken in this project is not suitable for a professional tool used by gynaecologists.

Chapter 7

Conclusion & Future Work

This project set out to explore various solutions to solve the problem of locating a vaginal pessary in 3-dimensional ultrasound data. Suitable algorithms were identified to determine the position of the ring crosscut. However, they can not yet reliably detect the orientation of the ring.

The problem was first subdivided into 3 sub-problems, each to be solved by a module. These 3 modules are the pre-processing module, the position analysis module and the rotation analysis module. After assigning each module specific targets it needed to fulfill, solutions were implemented and evaluated based on how well they achieved the given targets.

The final solution to the pre-processing module was found to be a concatenation of 3 filters: the image was first blurred using a vertical, 1-dimensional Gaussian filter. This output was then filtered by a LoG filter and finally a morphological opening filter was applied. This resulted in a binary image free of small spatial scale noise and showing the ring edges in high contrast to the background.

In approximately 70% of the cases the position can be determined accurately and reliably by analysing the contours of shapes in the filtered image. The orientation, however, can not be determined with sufficient confidence.

Possible solutions to determine the angle more precisely could lie in using the position estimate as a first guess for machine learning based approach e.g. an active shape model based approach. By using the location of the crosscut as a given point of the ring, the active shape would only need to optimize the ring model along 2 rotational axis.


Bibliography

- [1] M. Keisha A. Jones and M. Oz Harmanli, “Pessary use in pelvic organ prolapse and urinary incontinence,” [Online; accessed 27 Feb, 2020]. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2876320/pdf/RIOG003001_0003.pdf
- [2] F. Limbeek van den Noort, “Automatic segmentation of pelvic floor ring on 3d ultrasound,” [Online; accessed 19 Feb, 2020]. [Online]. Available: <https://www.ram.eemcs.utwente.nl/education/assignments/automatic-segmentation-pelvic-floor-ring-3d-ultrasound>
- [3] “Gynaecological imaging using 3d ultrasound (gynius),” [Online; accessed 23 Feb, 2020]. [Online]. Available: <https://www.nwo.nl/en/research-and-results/research-projects/i/45/28545.html>
- [4] M. Forouzanfar and H. Moghaddam, *Ultrasound Speckle Reduction in the Complex Wavelet Domain*, 01 2010, pp. 558–577.
- [5] Kanekin, “An introduction to ultrasonography,” Feb 2019, [Online; accessed 27 Feb, 2020]. [Online]. Available: https://www.idsmed.com/news/an-introduction-to-ultrasonography_403.html
- [6] W. F. Armstrong and T. Ryan, *Feigenbaums echocardiography*. Wolters Kluwer, 2019.
- [7] A. Fenster, G. Parraga, and J. Bax, “Three-dimensional ultrasound scanning,” *Interface Focus*, vol. 1, no. 4, 2011.
- [8] “Overview,” [Online; accessed 15 Feb, 2020]. [Online]. Available: <https://www.dicomstandard.org/about/>
- [9] D. Stipanich, “Introduction to digital image processing and analysis,” *Chronolab Vision, Chronolab A.G. Zug Switzerland*, 1994.
- [10] J. Rogowska, “Chapter 5 - overview and fundamentals of medical image segmentation,” in *Handbook of Medical Image Processing and Analysis (Second Edition)*, second edition ed., I. N. BANKMAN, Ed. Burlington: Academic Press, 2009, pp. 73 – 90. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780123739049500131>
- [11] F. Y. Shih, *Image processing and pattern recognition: fundamentals and techniques*. Wiley, 2010.
- [12] I. W. Hitoshi Matsumoto, Masato Ohtani, “Tree crown size estimated using image processing: A biodiversity index for sloping subtropical broad-leaved forests,” 2017, [Online; accessed 18 Feb, 2020]. [Online]. Available: https://www.researchgate.net/figure/Operations-of-mathematical-morphology-for-a-specific-structuring-element-Binary-images_fig4_318790302

Appendices

Appendix A

Global Threshold Filter




(a) Global Threshold Value = 50



(b) Global Threshold Value = 100



(c) Global Threshold Value = 150




(d) Global Threshold Value = 200

Figure A.1: Four levels of thresholding compared on a slice of patient ultrasound data

Appendix B

Adaptive Threshold Filter



(a) Reference Image



(b) Adaptive Threshold Block Size = 51



(c) Adaptive Threshold Block Size = 91




(d) Adaptive Threshold Block Size = 131


Figure B.1: Three adaptive thresholding block-sizes compared on a slice of patient ultrasound data. Neighboring pixel intensity weighted by a Gaussian distribution function

Appendix C


Sobel Filter




(a) Kernel Size = 1



(b) Kernel Size = 3



(c) Kernel Size = 5




(d) Kernel Size = 7


Figure C.1: Three sobel filter kernel sizes compared on a slice of patient ultrasound data

Appendix D


Derivative Filter Comparison




(a) Reference Image



(b) Scharr Filter



(c) Roberts Filter




(d) Prewitt Filter

Figure D.1: Derivative filters compared on a slice of patient ultrasound data


Appendix E

LoG Filter


E.1 LoG Filter Standard Deviation σ Variation




(a) LoG standard deviation $\sigma = 3$



(b) LoG standard deviation $\sigma = 4$




(c) LoG standard deviation $\sigma = 5$



(d) LoG standard deviation $\sigma = 6$

Figure E.1: Effect of the standard deviation parameter of the LoG filter on a slice of patient ultrasound data


E.2 LoG Filter after vertical Gaussian blur




(a) GaussVertical $_{\sigma} = 3$, LoG $_{\sigma} = 4$



(b) GaussVertical $_{\sigma} = 7$, LoG $_{\sigma} = 4$



(c) GaussVertical $_{\sigma} = 13$, LoG $_{\sigma} = 4$



(d) GaussVertical $_{\sigma} = 21$, LoG $_{\sigma} = 4$

Figure E.2: Effect of varying the standard deviation parameter of the vertical Gaussian blur filter GaussVertical $_{\sigma}$ before applying the LoG filter (LoG $_{\sigma} = 4$)

Appendix F

Morphological Operations



(a) LoG filtered image



(b) Kernel Size = 3



(c) Kernel Size = 4



(d) Kernel Size = 5


Figure F.1: Effect of kernel size on the binary morphological opening operation on the filtered ultrasound data (see figure E.2b)

Appendix G

Final Filter Algorithm On Ultrasound Images

The following images are represent a sample of the collection of experimentally obtained ultrasound scans of real patients.

G.1 Ultrasound Image 1




(a) Original slice of ultrasound data




(b) Filtered slice of ultrasound data

Figure G.1: Combined filter applied on Ultrasound Image 130448 at depth 112

G.2 Ultrasound Image 2




(a) Original slice of ultrasound data




(b) Filtered slice of ultrasound data

Figure G.2: Combined filter applied on Ultrasound Image 130185 at depth 109

G.3 Ultrasound Image 3



(a)



(b)

Figure G.3: Combined filter applied on Ultrasound Image 130208 at depth 110

G.4 Ultrasound Image 4




Figure G.4: Combined filter applied on Ultrasound Image 130328 at depth 110

Appendix H

Convolution Based Position Determination

H.1 Ultrasound Image 1




Figure H.1

H.2 Ultrasound Image 2




Figure H.2

H.3 Ultrasound Image 3




Figure H.3

H.4 Ultrasound Image 4





Figure H.4


Appendix I

Applied Contour Analysis

I.1 Ultrasound Image 1




(a) Contour of located ring edge result illustrated in the filtered 3D ultrasound slice




(b) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice

Figure I.1: Result of the contour based position analysis algorithm on ultrasound image 130448

I.2 Ultrasound Image 2




(a) Contour of located ring edge result illustrated in the filtered 3D ultrasound slice




(b) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice

Figure I.2: Result of the contour based position analysis algorithm on ultrasound image 130185

I.3 Ultrasound Image 3




(a) Contour of located ring edge result illustrated in the filtered 3D ultrasound slice




(b) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice

Figure I.3: Result of the contour based position analysis algorithm on ultrasound image 130208

I.4 Ultrasound Image 4



(a) Contour of located ring edge result illustrated in the filtered 3D ultrasound slice




(b) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice

Figure I.4: Result of the contour based position analysis algorithm on ultrasound image 130328


Appendix J

Correlation Based Angle Analysis


J.1 Ultrasound Image 130448



(a) Contour of located ring edge result illustrated in the filtered 3D ultrasound slice




(b) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice




(c) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice

Figure J.1: Result of the contour based position analysis algorithm on ultrasound image 130448


J.2 Ultrasound Image 130185



(a) Contour of located ring edge result illustrated in the filtered 3D ultrasound slice



(b) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice



(c) Contour of located ring edge result illustrated in the unfiltered 3D ultrasound slice

Figure J.2: Result of the contour based position analysis algorithm on ultrasound image 130185

J.3 Ultrasound Image 130208




Figure J.3: Result of the contour based position analysis algorithm on ultrasound image 130208

J.4 Ultrasound Image 130328

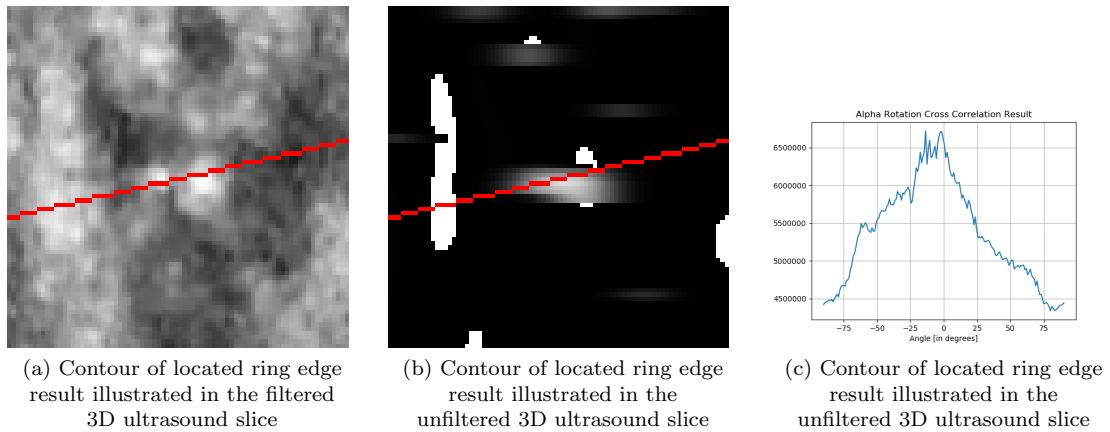



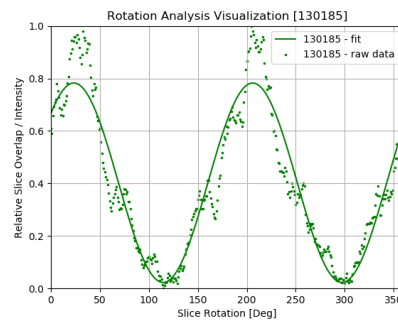
Figure J.4: Result of the contour based position analysis algorithm on ultrasound image 130328

Appendix K

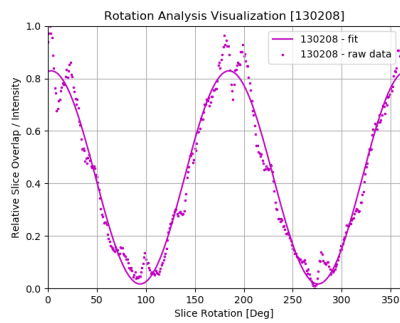
Slice Based Rotation Analysis




(a) Slice rotation analysis on unfiltered slices of ultrasound image 130448



(b) Slice rotation analysis on unfiltered slices of ultrasound image 130185



(c) Slice rotation analysis on unfiltered slices of ultrasound image 130208



(d) Slice rotation analysis on unfiltered slices of ultrasound image 130328

Figure K.1: Comparison of slice based rotation analysis on patient ultrasound images

Appendix L

Pre-Processing Analysis Results

Ultrasound Capture ID	Pre-Processing		
	Ring Crosscut Visible [Unfiltered]	Ring Visible [Filtered]	Success
130185	Yes	Yes	Success
130208	Yes	Yes	Success
130328	Yes	Yes	Success
130448	Yes	Yes	Success
130086	Yes	Yes	Success
130088	Yes	Yes	Success
130092	Yes	Yes	Success
130106	Yes	Yes	Success
130109	Yes	Yes	Success
130112	Yes	Yes	Success
130113	Yes	Yes	Success
130115	Yes	No	Failure
130121	Yes	Yes	Success
130122	No	No	Unsure
130130	Yes	Yes	Success
130142	No	No	Unsure
130148	Yes	No	Failure
130150	Yes	Yes	Success
130153	Yes	Yes	Success
130154	Yes	Yes	Success
130163	Yes	Yes	Success
130169	Yes	Yes	Success
130172	Yes	Yes	Success
130173	Yes	Yes	Success
130193	No	No	Unsure
130201	Yes	Yes	Success
130210	Yes	Yes	Success
130217	Yes	Yes	Success
130226	Yes	Yes	Success
130227	Yes	Yes	Success
130232	Yes	Yes	Success
130235	No	No	Unsure
130236	Yes	Yes	Success
130246	Yes	Yes	Success

Figure L.1: Pre-Processing Analysis Results

Appendix M

Position Analysis Results

Ultrasound Capture ID	Pre-Filtering Success	Position Determination Success
130185	Yes	Yes
130208	Yes	Yes
130328	Yes	Yes
130448	Yes	Yes
130086	Yes	Yes
130088	Yes	Yes
130092	Yes	No
130106	Yes	Yes
130109	Yes	No
130112	Yes	No
130113	Yes	No
130115	No	Unsure
130121	Yes	Yes
130122	Unsure	Unsure
130130	Yes	No
130142	Unsure	Unsure
130148	No	Unsure
130150	Yes	Yes
130153	Yes	No
130154	Yes	Yes
130163	Yes	No
130169	Yes	Yes
130172	Yes	Yes
130173	Yes	Yes
130193	Unsure	Unsure
130201	Yes	No
130210	Yes	Yes
130217	Yes	Yes
130226	Yes	Yes
130227	Yes	Yes
130232	Yes	No
130235	Unsure	Unsure
130236	Yes	Yes
130246	Yes	Yes

Figure M.1: Position Analysis Results

Appendix N

Rotation Analysis Results

Ultrasound Capture ID	Position Determination	Rotation Correlation Analysis		Slice Analysis		Manual Segmentation
	Success	Correlation Analysis Rotation	Correlation Analysis Rotation Error	Slice Analysis Rotation	Slice Analysis Rotation Error	Ring Rotation
130185	Yes	10	7.936864199	-46.105	64.0418642	17.9368642
130208	Yes	6	5.12263501	-6.616	7.49336499	0.8773649896
130328	Yes	-14	17.35895517	44.0045	40.64554483	3.35895517
130448	Yes	7	1.373478825	-17.702	23.32852118	5.626521175
130086	Yes	13	13	38.244	38.244	0
130088	Yes	48	45.36767211	4.7173	2.084972111	2.632327889
130092	No	-5		71.2		1.19717368
130106	Yes	-2	2	1.6984	1.6984	0
130109	No	74		-74.226213		26.71333577
130112	No	-77		-75.648		9.781997013
130113	No	-2		-31.659		0
130115	No	-5		-24.137		2.722970205
130121	Yes	-3	4.362254301	60.7776	59.4153457	1.362254301
130122	No	0		-30.75612		0
130130	No	-2		34.79		6.556673124
130142	No	-5		56.175		3.036585473
130148	No	40		38.8999		1.128787979
130150	Yes	-18	18.94067576	-46.548	47.48867576	0.9406757607
130153	No	66		65.9578		2.75457243
130154	Yes	1	1	10.292	10.292	0
130163	No	51		-47.4832		0
130169	Yes	4	4	28.7832	28.7832	0
130172	Yes	7	7	-18.123077	18.123077	0
130173	Yes	6	3.630593223	-27.99745	30.36685678	2.369406777
130193	No	11		-39.7625		2.272905538
130201	No	75		10.6595		1.013041587
130210	Yes	6	3.367675692	-61.933	64.56532431	2.632324308
130217	Yes	-2	3.067787625	54.56	53.49221237	1.067787625
130226	Yes	13	5.923921585	49.92085	42.84477158	7.076078415
130227	Yes	17	17	-1.0737	1.0737	0
130232	No	-82		53.5123		4.926923955
130235	No	0		33.0433		-0.7316716264
130236	Yes	-1	1	27.9884	27.9884	0
130246	Yes	10	10	-42.408	42.408	0

Figure N.1: Rotation Analysis Results