

Automatic Segmentation of Anterior Segment Optical Coherence Tomography Images

Thesis submitted in accordance with the requirements of the
University of Liverpool for the degree of Doctor in Philosophy by
Dominic Williams

January 2015



UNIVERSITY OF
LIVERPOOL

Abstract

Automatic segmentation of anterior segment optical coherence tomography (AS OCT) images provides an important tool to aid management of ocular diseases. Having precise details about the topography and thickness of an individual eye enables treatments to be tailored to a specific problem. OCT is an imaging technique that can be used to acquire volumetric data of the anterior segment of the human eye. Fast automatic segmentation of this data, which is not available, means clinically useful information can be obtained without the need for time consuming error-prone manual analysis of the images. This thesis presents newly developed automatic segmentation techniques of OCT images.

Segmentation of 2D OCT images is first performed. One of the main challenges segmenting 2D OCT images is the presence of regions of the image that generally have a low signal to noise ratio. This is overcome by the use of shape based terms. A number of different methods, such as level set, graph cut, and graph theory, are developed to do this. The segmentation techniques are validated by comparison to expert manual segmentation and previously published segmentation techniques. The best method, graph theory with shape, was able to achieve segmentation comparable to manual segmentation. Good agreement is found with manual segmentation for the best 2D segmentation method, graph theory with shape, achieving a Dice similarity coefficient of 0.96, which is comparable to inter-observer agreement. It performed significantly better than previously published techniques.

The 2D segmentation techniques are then extended to 3D segmentation of OCT images. The challenge here is motion artefact or poor alignment between each 2D images comprising the 3D images. Different segmentation strategies are investigated including direct segmentation by level set or graph cut approaches, and segmentation with registration. In particular the latter requires the introduction of a registration step to align multiple 2D images to produce a 3D representation to overcome the presence of involuntary motion artefacts. This method produces the best performance. In particular, it uses graph theory and dynamic programming to segment the anterior and posterior surfaces in individual 2D images with shape constraint. Genetic algorithms are then used to align 2D images to produce a full 3D representation of the anterior segment based on landmarks or geometric constraints. For the 3D segmentation, a data set of 17 eyes is used for validation. These have each been imaged twice so a repeatability measurement can be made. Good repeatability of results is demonstrated with the 3D alignment method. A mean difference of 1.77 pixels is found between the same surfaces of the repeated scans of the same eye.

Overall, a new automation method is developed that can produce maps of the anterior and posterior surfaces of the cornea from a 3D images of the anterior segment of a human eye. This will be a valuable tool that can be used for patient specific biomechanical modelling of the human eye.

Contents

1	Introduction	2
1.1	Overview	2
1.2	Aim and Objectives.....	3
1.2.1	Objectives.....	3
1.3	Unique Contribution of PhD	4
1.4	Background	4
1.4.1	The Cornea	4
1.4.2	OCT Data Sources	7
1.5	Specifications of Computer Used When Carrying Out Segmentation.....	10
1.6	Evaluation Criteria	10
1.6.1	2D Segmentation Evaluation	10
1.6.2	3D Evaluation	12
1.6.3	Statistical Techniques.....	13
1.7	Plan of Thesis.....	13
2	Literature Review.....	15
2.1	Introduction.....	15
2.2	Techniques for Imaging the Cornea.....	15
2.2.1	Optical Pachometry	15
2.2.2	Ultrasound Pachometry	16
2.2.3	Scheimpflug Photography.....	17
2.2.4	Optical Coherence Tomography	18
2.2.5	Anterior Segment OCT	21
2.2.6	Comparison of Different Methods	25
2.3	2D Image Segmentation	26
2.3.1	Pixel Based Methods.....	27
2.3.2	Active Contours	30
2.3.3	Graph Cut Segmentation	35
2.3.4	Graph Theory Based Methods.....	40
2.3.5	Relevance to this Work.....	45
2.4	3D Segmentation.....	45
2.4.1	3D Retinal OCT Imaging.....	46
2.5	Image Alignment Methods	48
2.5.1	Iterative Closest Point.....	48
2.5.2	Genetic Algorithms	48
2.5.3	Previous Use of Genetic Algorithms	52

2.5.4	Variations on GA.....	54
2.5.5	Using Zernike Polynomials to Model the Cornea.....	55
2.5.6	Relevance to Our Work.....	55
2.6	Conclusion.....	56
3	Image Segmentation in 2D.....	57
3.1	Introduction.....	57
3.2	Dealing with speckle noise in OCT image.....	58
3.3	Segmenting Images Using a Threshold.....	59
3.4	Common Segmentation Framework.....	60
3.5	Pre-Processing Steps.....	61
3.5.1	Removal of Iris and Central Noise Artefact.....	62
3.5.2	Rotation Error.....	63
3.5.3	Eyelid Removal.....	64
3.6	Level Set Method.....	67
3.7	Level Set with Shape.....	68
3.7.1	Initial Estimate of Segmentation.....	68
3.7.2	Evolution of Curve Using a Level Set Function with a Shape Constraint.....	69
3.7.3	Initial Results.....	74
3.7.4	Improvements to the Level Set Method using Dynamic Weighting....	75
3.7.5	Use of Texture Information to Improve the Model.....	80
3.7.6	Extending the Segmentation to Include the Iris.....	82
3.7.7	Alternative Ways to Generate Shape Energy.....	84
3.8	Graph Cut Segmentation.....	85
3.8.1	Improvements to Graph Cut Model.....	86
3.8.2	Using Boundary Term to Generate Shape.....	86
3.8.3	Using a Distance Function for Shape Term.....	88
3.9	Graph Theory Segmentation.....	88
3.9.1	Dijkstra Algorithm Method.....	89
3.9.2	Alternative Approach using Dynamic Programming.....	90
3.9.3	Detection of Additional Features.....	98
3.9.4	Using C++ to Improve Speed of Segmentation.....	101
3.10	Post Processing.....	101
3.11	Conclusion.....	102
4	3D Segmentation.....	104
4.1	Introduction.....	104

4.2	3D Level Set Segmentation	104
4.2.1	Segmentation Framework.....	104
4.3	3D Graph Cut Segmentation.....	108
4.3.1	Pre Segmentation Alignment	108
4.3.2	Creation of Shape Term	109
4.3.3	Segmentation	109
4.4	Graph Theory Segmentation in 3D	110
4.4.1	Segmentation	110
4.4.2	Registration	110
4.4.3	Genetic Algorithms	110
4.4.4	Restricting the Range of Solutions in GA.....	121
4.4.5	Comparison of Different Methods	122
4.5	Conclusion.....	122
5	Results	123
5.1	Introduction.....	123
5.2	Level Set based Segmentation	123
5.2.1	CVWS	123
5.2.2	Comparison of Three Different Weighting Techniques.....	126
5.2.3	Results of Texture Segmentation.....	128
5.2.4	Dynamic Weighting and Texture.....	130
5.2.5	Comparison of Different Level Set Methods.....	131
5.2.6	Discussion of Level Set Work	135
5.3	Graph cut.....	135
5.3.1	Binary Shape model	135
5.3.2	Distance Function Shape Model	138
5.3.3	Discussion of Graph Cut Segmentation	140
5.4	Graph Theory Segmentation.....	140
5.4.1	Dijkstra's Algorithm Method	140
5.4.2	Dynamic Programming Method	142
5.4.3	Graph Theory Time Comparison	144
5.5	Comparison of Different Segmentation 2D Techniques	145
5.6	3D Level Set.....	149
5.7	3D Graph Cut	152
5.8	3D Graph Theory.....	154
5.8.1	Non GA Registration Techniques	154
5.8.2	Genetic Algorithm Energy Function Variations.....	155

5.8.3	Genetic Algorithm Range Variations	157
5.9	Comparison of 3D Segmentation Techniques	158
6	Discussion and Conclusion	161
6.1	Discussion	161
6.2	Conclusion	163
6.3	Future Work	163
6.3.1	Comparison with Results from Other Imaging Modalities	163
6.3.2	Segmentation of SD OCT Images	164
6.3.3	Incorporation with Biomechanical Modelling	164
6.3.4	Improvements to the Segmentation Technique	164
7	References	165

List of Publications

Conference

Automatic segmentation of the cornea in optical coherence tomography images – MEDSIP 2012 04/07/2012 – presentation

A 3D segmentation framework for cornea segmentation in anterior segment OCT images using level set technique with shape prior, MIUA conference 2013 – presentation

Journal

D. Williams, Y. Zheng, F. Bao, and A. Elsheikh, "Automatic segmentation of anterior segment optical coherence tomography images," *Journal of Biomedical Optics*, vol. 18, pp. 056003-056003, 2013. [1]

D. Williams, Y. Zheng, F. Bao, and A. Elsheikh, "Fast segmentation of anterior segment optical coherence tomography images using graph cut", *Eye and Vision*, vol. 2, p. 1, 2015 [2]

D. Williams, Y. Zheng, P. Davey, F. Bao, M. Shen and A. Elsheikh, "3D Segmentation of Anterior Segment Optical Coherence Tomography Images Using Graph Theory and Genetic Algorithms", *submitted*

1 Introduction

1.1 Overview

Several eye healthcare therapies are based on the assumption that the eyes' topography and material stiffness vary consistently with, and can be predicted from, the age and medical history. In appreciation of the significant approximations made by this assumption, work is now progressing to optimise eye therapy for individual eyes [3]. The work is based on the use of numerical simulations that have been designed to closely represent the topography and material behaviour of an individual's eye, and predict the eye's response to mechanical actions such as surgery and impact. This project is intended to overcome an important obstacle related to the shaping of numerical simulations to an individual eye's topography as obtained clinically.

This work focuses on the analysis of optical coherence tomography (OCT) images of the cornea. The latest OCT instruments allow acquisition of 3-dimensional (3D) volumetric data and are capable of resolving the fine details of the anterior (front) part of the eye. In order to use these images for biomechanical modelling segmentation is needed to identify important boundaries in the image. Current algorithms supplied with anterior segment OCT machines are only able to accurately segment the central part of the images. Manual segmentation of the entire image is a time consuming process preventing fast analysis of images. The development of fully automatic segmentation techniques would allow clinically useful information to be obtained without the need for time consuming error-prone manual analysis of the images.

The main challenge faced when segmenting anterior segment (AS) OCT images is accurately detecting regions of low signal to noise ratio that occur either side of the centre of the image. In these regions the posterior surface between the cornea and the anterior segment chamber is not clearly delineated. In order to detect this region prior knowledge of the expected shape of the cornea must be used.

This thesis comprises of two main sections. The first develops fully automated image segmentation techniques capable of obtaining topographic information of the eye from 2D OCT images. The second section builds on this, developing segmentation and registration of a full 3D surface of the cornea.

The process of segmenting the images starts by using 2D OCT images of the anterior segment of the eye. Image segmentation is the task of splitting an image

into different regions. In this case the task is to split an image into two areas the object of interest, in this case the cornea and sclera, and background.

Segmentation of multiple regions of interest is also possible. There are various techniques that can be applied for image segmentation. Available techniques include finding a threshold to segment an image; a very simple and quick method that is vulnerable to noise. Using the vertical gradient of the image to locate edge points, segmentation based on gradient information can be carried out in a number of ways and may also be sensitive to noise. Active contours using a level set is an approach that uses various factors to segment an image including edge information, regional statistics and prior knowledge of the shape to be segmented. This technique has the advantage of being more robust to noise but is computationally expensive. Approaches taking a graph based approach to the problem including graph cut and graph theory based techniques have been tried. These work by representing images as a graph and require the development of energy functions that will correctly identify features on an image.

Following the segmentation of 2D images, the work is extended to obtain 3D topographical information. 3D segmentation can either be attempted by developing a fully 3D segmentation technique or by carrying out segmentation of multiple 2D images and combining them to produce a full 3D model. The different methods developed for use on 2D images are all extended to 3D. Those that are suited to it, such as level set and graph cut can be implemented in a fully 3D manner. Other techniques such as graph theory based segmentation cannot be easily extended to 3D, so 2D segmentation will be carried out on multiple images and then an alignment technique developed to produce a 3D surface from those images.

The 3D information to be acquired using OCT images could be used to produce patient specific numerical models of the human cornea, however this is beyond the scope of this thesis.

1.2 Aim and Objectives

The aim of the project is to develop a fully automated method that is able to create a 3D surface map that delineates the anterior and posterior boundaries of the human cornea from OCT images.

1.2.1 Objectives

In order to produce such a method there are two subsidiary steps that must be carried out

- Develop a fully automatic image segmentation program capable of segmenting 2D anterior segment OCT images
- Expand this method to utilise multiple 2D cross sections of the cornea to produce 3D image segmentation

1.3 Unique Contribution of PhD

As part of this PhD a number of novel techniques have been developed. There have been three new 2D segmentation techniques that have been developed. None of these involved the developed of new segmentation techniques but all involved the adaptation of existing techniques to the segmentation of anterior segment OCT images for the first time. The use of a shape based term to improve segmentation of anterior segment OCT images is novel. The first used a level set technique with a shape term. Use of level set based techniques on the segmentation of anterior segment OCT images has not been previously reported. The next technique developed was a graph cut segmentation using a shape term. This is the first time that graph cut segmentation with shape has been used on anterior segment OCT images. The final novel 2D segmentation technique was graph theory segmentation again with a shape term. This is the first time a shape term has been used on segmentation of anterior segment OCT images using this technique.

There has been no previous work that has attempted to validate 3D segmentation of anterior segment OCT images. The work carried out for this PhD is novel in this regard. The use of genetic algorithms in the way presented here presents a novel application of the technique to ocular OCT images.

1.4 Background

In order to carry out this segmentation, background knowledge of the structure of the cornea is necessary. This is presented in this section. Sources of OCT images are also described. All of the imaging was done by project collaborators since there was no suitable OCT imaging capability available at the University of Liverpool.

1.4.1 The Cornea

The human eye is a highly specialised organ that detects light from our surroundings and processes that information to the brain. Light is detected by specialised nerve cells in the retina called rods and cones. This is an important part of the eye; the purpose of other structures of the eye is to focus light onto the retina and to provide the needed support for the eye's internal components.

The general shape of the eye is approximately spherical in nature. Its outer tunic consists of two parts, a smaller anterior segment and the larger sclera. The cornea is located at the front of the eye and forms a major part of the anterior segment. Together with the sclera, the cornea provides a tough protective envelope which protects the important ocular tissues. The cornea has a higher curvature than that of the sclera and forms about one sixth of the total surface of the eye.

The surface of the cornea is where most of the refraction in the eye occurs; there is a large difference in refractive index between air and corneal tissue that changes little with age. The transparency of the cornea is one of its most important properties, although it must also provide a tough physical boundary against trauma and infection [4]. The cornea is composed of five layers, these are the corneal epithelium, Bowman's layer, the corneal stroma, Descemet's membrane and the cornea endothelium [5]. The cornea has no blood vessels and has a highly regularised arrangement of fibrous, cellular and extracellular components: these ensure the transparency of the cornea.




Figure 1: Labelled diagram of the human eye taken from [6].

1.4.1.1 **Corneal Epithelium**

The corneal epithelium is the outermost layer of the cornea. It is 50-60 μm thick and consists of 5 or 6 layers of flattened, nucleated and nonkeratinised cells. The anterior surface has numerous microvilli and microplane which act to stabilise the pre-corneal tear film. New cells grow in the limbal basal cell layer and displace

existing cells superficially and centripetally. Cells slide to rapidly repair damage to the integrity of this layer.

1.4.1.2 Anterior Limiting Lamina

The anterior limiting lamina, or Bowman's layer, is 8 – 12 μm thick and consists of fine, randomly arranged collagen fibrils. It connects the epithelium with the stroma. The boundary with the epithelium is well defined, while the posterior boundary merges with the stroma.

1.4.1.3 Corneal Stroma

This is the region that takes up the vast majority of the cornea and is generally between 400-500 μm thick. It is made of dense connective tissue containing 2 μm thick layers of collagenous lamellae, most of which are orientated parallel to the surface [7]. Keratocytes are found between these layers, these are extremely flattened fibroblasts, which act to hold the collagen bundles in place and may help with diffusion of metabolites. This region normally contains no blood or lymphatic vessels.

1.4.1.4 Posterior Limiting Lamina

This is a thin, homogenous layer between the stroma and endothelium 8 – 12 μm thick. There are two parts to it an anterior banded section and a nonbanded posterior section. Microscopic deposits appear on the edge of the region with increasing age. If this layer is disrupted it tends to curl towards the anterior chamber.

1.4.1.5 **Corneal Endothelium**

The corneal endothelium is responsible for regulating the supply of fluid to the rest of the cornea. The cornea loses fluid constantly through evaporation causing the thickness to decrease as the day progresses. The cells in the layer are 5-6 μm in height and 18-20 μm in diameter and form an uninterrupted polygonal array. They have a very low regenerative capacity and cell density fall with age from 4000 cells/ mm^2 at birth to 2000 cells/ mm^2 in old age. Damage to cells and low density causes oedema and swelling of the stroma resulting in a loss of transparency.




Figure 2: Labelled cross section of the human cornea from [8].

1.4.1.6 **Boundaries of Interest**

This study aims to detect the anterior and posterior boundaries of the cornea using OCT images. The anterior boundary corresponds to the air-epithelium interface. Bowman's membrane is visible in some OCT images, but not all of them, and not for the entire width of the cornea. For this reason I do not attempt to segment this layer. The posterior boundary corresponds to the location of the endothelium.

1.4.2 **OCT Data Sources**

There was no capability to carry out anterior segment OCT imaging in the University of Liverpool. All the imaging was carried out by collaborators. Each set of image acquired will be discussed in turn.

1.4.2.1 **Initial Data Set**

In order to develop an image segmentation program for the anterior segment, OCT images were needed. The images initially used were collected at the Singapore National Eye Centre. The Visante anterior segment optical coherence tomography

(AS OCT) system manufactured by Zeiss was used. This is a time domain system that uses $1300nm$ infrared light to obtain cross sectional images of the anterior segment with a scanning rate of 2000 axial scan per second. The image contains 256 A-scans in $16mm$ with 1024 points per A-scan to a depth of $8mm$. The images have a transverse resolution of $60\mu m$ and an axial resolution of $18\mu m$. The images were outputted by the machine as 816×636 pixel JPEG files. Informed consent was obtained from all participants in the study. Figure 3 shows an example image obtained using this machine.




Figure 3: Example 2D anterior segment OCT image.

1.4.2.2 **Second 2D Data Set**

In order to validate the automatic segmentation methods developed, comparison against manual segmentation must take place. A second data set was acquired for this purpose. 39 anterior segment OCT B scan images through the centre of the cornea from healthy eyes (one per subject) were acquired by the Visante AS OCT system (Carl Zeiss Meditec, Dublin, CA) in Wenzhou Medical University, China for the purpose of evaluation in the study. This is the same system as the initial data set. The images were corrected for refractive index using the built-in software of the system; this correction is unlikely to affect the automatic segmentation results. The anterior and posterior boundaries of all images were later segmented manually by two expert ophthalmologists.

1.4.2.3 **3D Data Sets**

All the data used for 2D segmentation came from time-domain (TD) OCT devices. This is because they are currently the most widely available commercially available systems capable of imaging the full width and depth of the cornea. Previous studies

have been carried out using spectral-domain (SD) OCT to investigate the cornea, but these have mainly involved the use of custom built machines. One of the advantages of SD OCT over TD OCT is the increased speed of measurements. This is very useful in obtaining 3D data. The majority of OCT machines currently available are set up to take 2D images slicing into the eye. Multiple 2D slices of the cornea can be combined to produce a 3D image.

Initial 3D work was carried out on a single set of images of an eye obtained through SD OCT containing 32 individual 2D images of the cornea. These data were acquired using a custom built SD OCT system in Wenzhou Medical University, China. Synthetic data were created to supplement this data set. These images were created simulated the cornea as the area between two ellipsoids, speckle noise was added to this data and part of the cornea removed to simulate low signal to noise regions present in real data.

Due to the roughly ellipsoidal structure of the cornea, a series of 32 cross-sectional, radial images of the cornea all across the centre of the cornea but rotated relative to each other are taken to represent the cornea. Figure 4(a) demonstrates the scanning pattern and Figure 4(b) illustrates the scan in the horizontal direction.

Another source of 3D data was through collaboration with Dr Pinakin Davey at Western University of Health Sciences, USA. An initial set of 8 eyes each with 16 B scans again in a radial pattern was procured. These were taken using the same Visante OCT system that was used to acquire the 2D data.

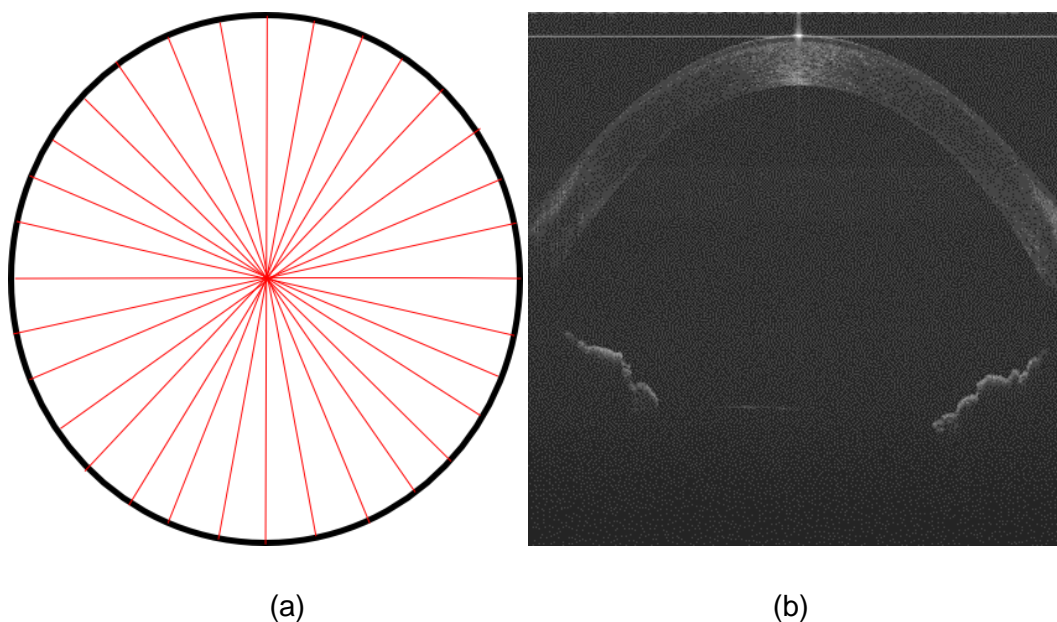


Figure 4: Illustration of the 3D AS-OCT scanning pattern. (a) Diagram showing layout of radial scans. The black circle represents the cornea and each red line is a B scan of the cornea. Note that only 16 scans are shown here for clarity. (b) An example scan in the horizontal direction. For the purpose of demonstration, the brightness and contrast have been adjusted.

1.4.2.4 Repeated Measurements

In order to carry out validation of our 3D segmentation and alignment technique a repeatability test was carried out. Additional data for this test was acquired from the Western University. The same Visante OCT machine was used to acquire 16 images of an eye arranged in a radial pattern. Each eye was imaged twice by this method to enable a repeatability study to be carried out. Images of a total of 17 eyes were acquired for this study. These came from 11 different people, 5 people had images taken of a single eye and 6 people had images taken of both eyes.

1.5 Specifications of Computer Used When Carrying Out Segmentation

The different methods developed here have been tested for speed and accuracy. In order for effective comparison of the time taken for segmentation it is important to know the specifications of the machine that was used for the segmentation. All segmentation was carried out using a Win7 PC with Intel Core i5-2320 CPU @3.00GHz and 4.00 GB RAM.

1.6 Evaluation Criteria

1.6.1 2D Segmentation Evaluation

In order to evaluate the success of segmentation techniques it is necessary to compare the results of automated segmentation techniques with those done by expert manual observers. There are numerous comparison techniques: this study uses three measures to evaluate success; Dice's similarity coefficient (DSC), mean unsigned surface positioning error, and the Hausdorff distance (HD).

1.6.1.1 Dice Similarity Coefficient

The DSC is an area similarity method defined by the following expression

$$DSC = \frac{2|X \cap Y|}{|X| + |Y|} \quad (1)$$

where X and Y are the two segmentations to be compared, in this case the manual and automated segmentation results. DSC has a range between 0 and 1. The higher the DSC value, the more similar the two segmented regions are. This was first developed by Dice when looking at ecological association of species [9] but has

been used widely for evaluating images since then [10]. Figure 5 is a diagram that shows how this method works.




Figure 5: Diagram showing Dice Similarity Coefficient (DSC). The two regions to be compared are X outlined in red and Y outlined in green. The area of overlap is yellow. DSC is defined as the area in yellow divided by the total area that X and Y encompass.

1.6.1.2 Mean Unsigned Surface Positioning Error

The mean unsigned surface positioning error (MSPE) is a method for quantifying the similarity between two lines. It differs from the DSC in that it looks at individual lines rather than the area as a whole. The unsigned surface positioning error is defined by

$$MSPE = \frac{\sum_{i=1}^n |z_a^i - z_m^i|}{n} \quad (2)$$

where z_a^i is the z value of the automated segmentation at point i , z_m^i is the z value of the manual segmentation at point i and n is the number of points. The method has been widely used for evaluation of image segmentation techniques [10]. Figure 6 is a diagram that shows how this method works.




Figure 6: Diagram of mean unsigned surface positioning errors (MSPE). The red and green lines are to be compared. The mean unsigned distance corresponds to the black lines with measures made at set points along the line. The distance is unsigned so it does not matter which line is on top

1.6.1.3 Hausdorff Distance

The mean 95% Hausdorff distance [11] is a more stringent measure that compares the difference between the two boundaries. The Hausdorff distance from set A to set B is defined as

$$HD(A, B) = \max_{a \in A} (\min_{b \in B} (|a - b|)) \quad (3)$$

where A and B are sets of boundary points from the two images to be compared, a and b are the corresponding individual points within each set. The measure is in effect a measure of the distance of the worst fitting point from the data set. In order to reduce the effect of outliers on the data set, a modification of the Hausdorff distance can be used. Here the 5% largest distances were removed. Then the maximum of HD(A,B) and HD(B,A) was taken for each image [12]. Perfect alignment is represented by a Hausdorff distance of 0.

1.6.2 3D Evaluation

For the 3D methods, two sets of data were acquired to enable repeatability tests to be carried out. Repeatability was chosen to evaluate the success of the techniques since it is not practical to carry out manual registration. Two measures were used to assess repeatability. The first was a simple height difference comparison. The height difference between corresponding points on each of the two surfaces is calculated across the surface. An unsigned mean was then taken of this over the entire area. This measure is defined by

$$R = \frac{\sum_i (z_i^a - z_i^b)}{n} \quad (4)$$

where z_i^a is the z coordinate of the ith point on surface a, z_i^b is the z coordinate of the ith point on surface b, i is valued from 1 to n and n is the number of points present on both surfaces being compared.

The other measure used was a coefficient of repeatability (CoR) this is a measure of the repeatability of multiple measurements of the same things. Measurements of the corneal thickness were used to evaluate this. The CoR is defined by

$$R = \frac{2 * sd(Ta - Tb)}{mean(Ta - Tb)} \quad (5)$$

Where sd is the standard deviation of the difference between the two measurements Ta and Tb .

For both measures of repeatability a lower value indicates a better match.

1.6.3 Statistical Techniques

In order to test if differences between methods are statically significant a number of techniques were used. The statistical tests were all calculated using the software package SPSS Statistics for Windows (version 20, IBM, Armonk, NY, USA).

When comparing more than two techniques an ANOVA test was carried out. This test enables quick comparisons between different techniques to see if differences between them are significant. Post hoc analysis can then be carried out to assess the significance of specific differences between methods.

When two methods were being compared directly, a paired t-test is more appropriate to use. The two methods are compared over the same set of images. Using a paired t-test will account for differences in results that are due to differences between the images rather than the alignment techniques.

1.7 Plan of Thesis

The remainder of the thesis is split into 5 chapters.

Chapter 2 presents a literature review. Previous work investigating different imaging techniques for the cornea is discussed. Different image segmentation techniques, as well as different registration techniques, are then discussed.

Chapter 3 describes the 2D segmentation techniques developed. Initially, existing segmentation techniques that have been used on the anterior segment are described. Following this a number of techniques that have been developed as part of this study are explained. The first group of methods introduced are based on level set techniques. Methods based on graph cut techniques are then introduced and finally graph theory segmentation is explored. For each group of methods a number of novel techniques have been developed specifically aimed at segmenting anterior segment OCT images.

Chapter 4 describes 3D segmentation and alignment. There has been no previous work compatible with the imaging system used in this thesis. Three different approaches for developing novel techniques are discussed. These are based on the 2D segmentation methods developed in previous chapter. Techniques to align sets of 2D images to produce 3D maps are also described.

Chapter 5 contains the results of the segmentation methods. 2D automated segmentation programs developed here are compared to manual segmentation.

The results of the different techniques are presented in the same order that they were described. Different variations of the techniques are compared within each broad category of method. Then one of each type of method is compared so the best method can be decided upon. The second half of this chapter discusses the results of the 3D segmentation and alignment techniques. The results of a repeatability test are presented and used to assess the performance of each method.

Chapter 6 contains a discussion of future steps that can be carried out based on this work. It also contains a summary of the main findings to conclude the work.

2 Literature Review

2.1 Introduction

This chapter will first review different techniques that have been used for the purpose of imaging the cornea. The advantages and disadvantages of techniques are discussed. This review will show why OCT is a promising technique for use in the project. We will then go on to discuss image segmentation techniques with special focus being placed on those that have previously been used on ocular OCT images. 2D segmentation techniques are discussed first followed by 3D segmentation. Image alignment problems inherent in 3D segmentation are discussed next. Finally a concluding section will present important findings and point at areas of investigation for the current study.

2.2 Techniques for Imaging the Cornea

Given the cornea's importance in human vision it has been subject to much investigation. There are several methods that can be used to measure corneal thickness [13]. These include ultrasound pachometry, optical pachometry, Scheimpflug photography and OCT, as well as techniques that combine different methods. Each method will be discussed separately and comparisons will be made.

2.2.1 Optical Pachometry

2.2.1.1 *Slit Lamp Based Method*

Optical pachometry is based upon the use of a slit-lamp. The technique requires a trained observer to judge where the light is reflected from within the cornea. The location of one edge is ascertained and then the relative distance to the other edge can be measured using a sliding scale. This technique can measure central corneal thickness to within $6\mu m$. The use of a set of fixation lights helps to assist in the alignment of the slit lamp [14].

One of the major disadvantages of this method is that it is a subjective method relying on the observers' judgement about location of boundaries [15]. It also suffers from potential inaccuracies due to possible errors due to alignment and non-uniform specula reflection. This technique was largely replaced by ultrasound methods by the late 1980s.

2.2.1.2 *Clinical Specula Microscope*

This is a variation on slit lamp optical pachometry. A slit lamp is used to illuminate the cornea at an angle. The light from this is reflected from both the posterior and anterior surfaces of the cornea. The light reflected from both surfaces is

approximately parallel so can be observed with a microscope. In order to improve the accuracy of the results an aperture is placed in front of the objective lens. This reduces the range of possible deviations from the correct angle [16].

Comparisons between this method and ultrasound pachometry found significantly lower values of corneal thickness when the specula microscopy method was used (32 μm thinner) [17]. Lower inter observer variation was observed in the specula method, indicating it could be more useful for long term studies [18]. This is probably due to the fact that the position of the cornea is fixed using a guide light and no manual placing of an ultrasound probe needs to take place.

2.2.1.3 Orbiscan System

This is an optical system that can be used to measure the topography of the human eye. It measures the anterior and posterior corneal elevation, surface curvature, and corneal thickness. Two slit lights are used to illuminate the human eye at a 45 degree angle. Slits are projected sequentially onto the eye from both left and right. Information on where this light is reflected from is then used to calculate axial curvature, elevation of anterior and posterior surface, and corneal thickness in a central region 8 to 10 mm in diameter. Pachometry data is determined by the difference in elevation between the anterior and posterior surface of the cornea [19].

Comparisons between this system and ultrasound pachometry techniques have found similar accuracy and precision in both systems. Orbiscan measurements were found to be 23-28 μm greater than ultrasound measurements of the same eye [20]. This means that this system cannot be used interchangeably with ultrasound methods. This discrepancy could be due to the fact that ultrasound methods are contact whereas the Orbiscan system is a non contact method. It could also be due to the Orbiscan system including the hydrated mucus gel that covers the corneal surface and has a thickness of up to 40 μm , as part of the cornea [21].

2.2.2 Ultrasound Pachometry

Ultrasound pachometry is a widely used technique which has been around for over 25 years and is frequently used in the clinical environment [22]. Ultrasound pachometry works by emitting high frequency pulses of sound which are reflected by the anterior and posterior surfaces in the cornea. These reflected pulses can be measured and the time difference between pulse emission and signal can be used to determine the location of the edges. A transducer probe which emits high frequency sound pulses must be placed on the cornea to enable measurements to be made. Patients must receive a topical anaesthetic.

The method has several advantages. The technique requires minimal observer judgement making measurements more reliable and consistent. It is cheaper, faster and more portable than optical pachometry, a previously used technique. One disadvantage of this method is that it requires contact between the ultrasound probe and the cornea. This introduces potential errors since if too much pressure is applied with the tip, incorrect readings will result. In order to find the centre of the cornea patients are asked to stare at a fixed spot and the centre of the pupil is estimated to be the centre of the cornea. No fixation lights are used to control patient gaze. One study [23] found a high degree of reproducibility of ultrasound pachometry results for both inter observer and intersession measurements. This shows that the issues mentioned above do not present a fundamental problem for the technique.

Ultrasound pachometry is a single point measurement technique and doesn't produce a map of the cornea. Ultrasound biomicroscopy (UBM) is a variation on this technique that is able to produce maps of the cornea [24].

2.2.2.1 *Ultrasound Biomicroscopy*

UBM requires the use of a water bath to achieve contact between the eye and the probe. Early models required patients to lie down to take measurements. More recently the Paradigm UBM model (Paradigm Medical Industries Inc., Salt Lake City, Utah, USA) uses a handheld device that can measure patients while sitting [25]. UBM has been used to detect abnormalities in anterior segment angle and ciliary body. Currently it is not used as a standard diagnostic tool but has proved a useful tool for research purposes.

2.2.3 *Scheimpflug Photography*

This is a variation of the slit lamp method of imaging the anterior segment. Light is directed on the eye and a series of images are taken at an angle to each other. These are combined to generate cross-sectional images of the anterior segment of the eye. These give information about where light is reflected from the eye and so can be used to determine the density of the cornea or lens [26]. When using the Scheimpflug method to look at the lens it produces similar results to slit lamp based classification but is less dependent on the skills of the person performing the examination. When compared to techniques like OCT or ultrasound biomicroscopy (UBM), Scheimpflug photography is superior in respect of quantification of light scattering in the cornea and lens. However it cannot be used to visualise the complete anterior chamber angle which OCT and UBM are capable of doing.

2.2.4 Optical Coherence Tomography

Optical coherence tomography (OCT) is a widely used medical imaging technique. Since it was first developed in 1991 [27], it has been rapidly developed as a non invasive optical medical diagnostic method. It generates cross sectional images by measuring time delay in backscattered light. OCT was first used for in vivo studies of the human eye, specifically looking at the retina. The transparent nature of ocular media, where only minimum optical attenuation and scattering occurs, provides easy access to the human retina. This has meant that ophthalmic diagnosis is one of the most clinically developed OCT applications. For example around 50% of all OCT publications up to 2008 were in ophthalmic journals.

OCT of the anterior segment of the eye was first demonstrated in 1994 [28]. There have been many different uses of anterior segment OCT [29]. These include evaluating anatomical outcomes of refractive surgery, imaging of cataracts, anterior chamber angle assessment, monitoring corneal oedema, and routine clinical examination of the anterior segment. However, the extra expense of OCT machines compared to ultrasound equipment and the established nature of ultrasound measurements has meant that ultrasound continues to be used as the standard technique for measuring central corneal thickness.

Using ultrahigh resolution OCT, an axial resolution of 2 – 3 μm can be achieved, and standard OCT systems have a maximum resolution of 10-15 μm [30], better than can be achieved using ultrasound. The technique works in a similar way to ultrasound but detects back scattered light instead of sound.

This method has several advantages over ultrasound methods. It has a higher resolution than conventional ultrasound due to shorter wavelength of light, the method is non invasive and non contact, and it can give quantitative morphological information [31]. OCT has been used to look at the cornea. Measurements of central corneal thickness can be made using this technique. This is done using an automatic algorithm to isolate the cornea in the central region. Studies have good repeatability indicating that OCT is a suitable technique to make these measurements.

The technology behind OCT has developed and improvements in both scanning speed and resolution have been achieved. There are two types of OCT systems, time domain and spectral domain OCT. Time domain OCT will be discussed first, followed by spectral domain OCT.

2.2.4.1 Time Domain Optical Coherence Tomography

OCT works by measuring the time delay of backscattered light incident on a sample. The time delay of the backscattered light cannot be measured directly due to the high speed of light, so instead an interferometer is used. The light emitted from a source is split into two paths using a beam splitter, known as the reference and sample arms.

Light in the reference arm travels to a mirror at a known distance, it is then recombined with the light from the sample arm, which has been back scattered by the sample. The reference mirror is moved while the sample is kept stationary, creating an interference pattern is measured by the detector. The coherence length of the light source limits the depth of possible measurements. A measurement of reflectivity against depth can then be obtained from the detector; this single scan of reflectivity against depth is referred to as an A scan. A series of these A scans taken along a line is known as a B scan and can be used to create a 2D image of the subject.




Figure 7: Schematic diagram of Time Domain Optical Coherence Tomography taken from [32] and modified. SLD is the light source which can be a super luminescent diode, other light sources can also be used. For time domain OCT the sample is held still and the reference arm is moved to generate an interference pattern after light is recombined and sent to detector.

The axial resolution of an OCT image is primarily determined by the coherence length of the light source. This is determined by the following equation [33]

$$l_c = \frac{2 \ln(2)}{\pi} \frac{\lambda_0^2}{n \Delta \lambda} \quad (6)$$

where λ is the central wavelength of the light source in question, $\Delta\lambda$ is the bandwidth of the light source and n the refractive index of the sample. The choice of wavelength is therefore important in OCT; whilst a shorter wavelength is desirable in that it would give an improved resolution, tissue absorption must also be considered. Haemoglobin absorption dominates at wavelengths shorter than 600nm and water absorption becomes important above 1000nm [34]. This means that the standard wavelengths used are in the 800-900nm region. More recently longer wavelength light sources have also been used with a wavelength of 1310nm being used in several systems. The use of infrared wavelengths has the advantage of not dazzling the subject. Increased absorption at longer wavelengths also allows for a higher power of light to be used, since absorption reduces retinal exposure.

2.2.4.2 Spectral Domain OCT

A more recent development in OCT is spectral domain (SD) or Fourier domain OCT. In this system the reference mirror is kept stationary and the OCT signal is acquired as a function of frequency using a spectrometer as detector [35]. The resultant spectrum can be Fourier transformed to determine where the backscattering of the light occurred. SD OCT can be swept source where the emission wavelength is tuned rapidly over a broad wavelength, in this case the set up is the same as TD OCT and a single photodetector can still be used. The other alternative is to use a broadband light source and use a spectrometer to record the signal.




Figure 8: Schematic Diagram of Spectral Domain OCT. The detector in this case must be a spectrometer [32]. Here SLD is a light source. For SD OCT the light source must either be broadband source or a have a varying wavelength light emitted. The reference mirror is held still in the device.

One of the main advantages of SD OCT over TD OCT is the increasing scanning speed available. Since SD OCT contains no moving parts a much higher acquisition speed is achieved. One problem faced when making optical measurement is involuntary eye movements by patients. Improving acquisition speed reduces the effect of this problem. In 2008 the fastest commercial TD OCT system acquired 2,000 A-scans per second; in contrast a SD OCT system then available could achieve 26,000 A-scans per second or higher [36]. This opens up the possibility of creating 3D scans of the eye. It also allows for a higher density of A scans, thereby increasing sensitivity and reducing speckle size and so noise in the images.

2.2.5 Anterior Segment OCT

Whilst OCT has been used most widely to study the retina it has also been used extensively to study the cornea as well. Accessing the Scopus online database in 2012 [37] the search term 'optical coherence tomography AND cornea' revealed 1,109 results compared to 7,823 results using the search term 'optical coherence tomography AND retina'.

When looking at the anterior segment a longer wavelength can be used. Longer wavelength light has the advantage of increasing the penetration depth. A higher power of light can also be used since water absorption at this frequency acts to

shield the retina from receiving too much light [38]. Scattering is reduced at longer wavelengths which leads to an improved penetration depth when 1310 nm is used instead of 803 nm [39].

Work has been carried out developing SD OCT systems with a larger scanning depth. One group investigating accommodation designed a custom built OCT system that was capable of achieving a scanning depth of 11.87mm [40]. The scanning speed of this system was 4,000 A scans per second. This is faster than TD OCT systems and able to achieve similar scan depths to those they can achieve. Several different uses of anterior segment (AS) OCT will be discussed in the following sections.

2.2.5.1 Anterior Chamber (AC) Angle Assessment

Glaucoma is one of the most common causes of blindness and irreversible visual loss. It is caused by poor filtration of fluid in the eyeball and, if untreated, results in higher intraocular pressure which can cause permanent nerve damage and so blindness. Intraocular fluid is usually filtered through the anterior chamber angle. If this angle becomes too small then angle closure glaucoma occurs. It is estimated that this will cause 5.8 million people worldwide to go blind by 2020 [41]. This means that screening for glaucoma is very important especially since it is asymptomatic in the early stages, resulting in cases not being diagnosed till the disease is advanced and has affected the patient's vision. Measuring the anterior chamber angle is used to detect angle closure glaucoma and decide the treatment method.

There has been a considerable amount of research into using AS OCT for anterior chamber angle assessment. The use of light allows a higher spatial resolution than ultrasound to be achieved. The non-contact nature of OCT permits more comfortable examinations and removes possible mechanical distortion of the angle.

In 2005 the commercially available Zeiss Visante OCT was approved for AC angle imaging. This uses 1300nm wavelength and was able to measure 2,000 A scans per second with an axial resolution of 10-20 micrometers. Several studies were carried out based on manual identification of the sclera spur from OCT images. This was then used to determine the AC angle and the sclera spur location, which could be determined in 72% of cases [42].

Advances in OCT technology and the introduction of SD OCT meant that higher resolution images can be produced in a shorter time. High definition OCT with

improved axial and transverse resolution promises higher quality images and more reliable angle assessment. However, the shorter wavelengths used in HD OCT devices, 840nm as opposed to 1310nm, mean that penetration depth is reduced. As a result of this the sclera spur, previously used for angle identification, is not well defined. However, Schwalbe's line (at the back of the cornea) is well defined and can be used as an alternative method of angle assessment. Attempts have been made to develop a fully automated system for detecting this line [43].

An alternative SD OCT system that is capable of imaging the entire anterior chamber has is the Casia SS-1000 OCT (Tomey, Nagoya, Japan). This is a commercially available swept source OCT system using a wavelength of 1310 nm [44]. This machine comes with automated software that is able to detect posterior and anterior surfaces of the cornea. With the manual identification of a few landmark points it can be used to measure the anterior chamber angle. Low variability of measurements of the anterior chamber angle has been shown using this machine. One group found the inter class correlation coefficient (ICC) was >0.83 [44]. The segmentation technique used on this machine is unknown. The software is proprietary and has not been published anywhere. This means it is not possible to compare the segmentation technique to techniques developed by others.

2.2.5.2 Anterior Chamber Biometry

Measurements of anterior chamber dimensions are important for many different purposes, including the sizing of ophthalmic devices such as intraocular lenses (IOL) [45], corneal implants, microkeratomes and contact lenses. Both UBM and AS OCT are capable of measuring the anterior chamber. The non contact nature of OCT is a major advantage. Studies using AS OCT to measure anterior chamber have found good repeatability in their results [46], showing that it is a good technique. When carrying out these measurements manual identification of key points in the image must be carried out.

2.2.5.3 Pathologies of Anterior Segment

Anterior segment OCT has been used to investigate various different pathologies of the anterior segment. One study looked at patients with micro bacterial keratitis, which is an infection associated with contact lens wear, trauma, and ocular surface disease, and can lead to visual loss. Current diagnosis is based on slit lamp examination but AS OCT has been shown to be able to provide an objective assessment of the disease [47].

Work has also been carried out investigating anterior segment tumours using OCT. Here problems were encountered when looking at larger tumours, as OCT was unable to penetrate the entire tumour due to increased pigmentation [48], when looking at smaller or hypopigmented iris tumours this problem did not occur. OCT was also able to differentiate between solid and cystic lesions in the iris [39].

The noncontact nature of AS OCT makes it ideal for measurements of fragile eyes. Trauma classification is very important to evaluate severity and plan a management course. AS OCT has been shown to be able to detect displacement of the crystalline lens and occult sclera perforation [49] in a patient suffering loss of vision after blunt trauma from a clenched fist. Using ultrasound in this case would be unsuitable since a method involving contact would risk further injury to the eye and OCT can reveal a lot more than slit lamp examination.

2.2.5.4 *Laser in Situ Keratomileusis Flap Thickness*

Laser in situ keratomileusis (LASIK) is a commonly carried out corneal refractive procedure. Following surgery, monitoring LASIK flap thickness is very important to ensure adequate recovery. AS OCT has been used to carry out these measurements, as have ultrasound techniques. Studies have found good repeatability of measurements using OCT [50]. Studies have also been carried out comparing results from spectral domain and time domain OCT systems. These found that there was no significant difference in mean flap thickness measurements between the two techniques. One study found better intra observer agreement when using SD OCT [51].

2.2.5.5 *Keratoconus Treatment*

Keratoconus is a bilateral non inflammatory progressive corneal disease characterised by changes in corneal collagen structure, decreased rigidity of the cornea, and corneal thinning. Patients will suffer from progressive corneal deformation leading to decreased vision if untreated. Treatment options include using contact lenses or glasses to correct for refractive error. Patients who are intolerant to contact lenses can be treated through intrastromal corneal ring implantation. OCT can be used to measure the placement depth and identify patients at risk of depth related complications [52].

Another newer treatment method is corneal collagen crosslinking. Eye drops and UV light are used to form crosslinks between corneal collagen fibres which act to increase corneal rigidity. Following on from the treatment a demarcation line appears that separates the treated area from the rest of the corneal stroma. AS

OCT has been demonstrated to be able to detect the line and accurately measure its depth, which is thought to be correlated to the effectiveness of the treatment [53].

2.2.5.6 Other Uses

Anterior segment OCT has been shown to be able to accurately image the anterior segment of the eye. This has led to it having research roles in other corneal surgery applications, including monitoring Phakic Intraocular lenses, corneal transplantation, cataract surgery, glaucoma surgery and other surgery [39, 54-56].

2.2.6 Comparison of Different Methods

Measuring corneal thickness is something that is useful to do over a long time period. It is hard to get absolute measurements for in vivo readings using the various techniques since measuring an eye in vivo and then dissecting it to measure actual thickness is not possible. The fact that the different techniques show different thicknesses when looking at the same eye is therefore a strong disincentive to change technique. In evaluating a method for its suitability for measuring central corneal thickness repeatability is the criterion used. A technique with good repeatability allows measurements taken years apart to be effectively compared. Ultrasound methods have proved to be just as good as other techniques developed in this regard and so have continued to be the most widely used.

Other methods such as the Orbscan technique and OCT are able to give topographical information about the cornea as opposed to just giving information on corneal thickness. This makes them more useful in investigating patients known to have more complex problems.

A comparison of ultra sound pachometry, OCT, and Pentacam (a Scheimpflug system) showed that they achieve different results for central and peripheral corneal thickness [57]. The authors concluded that this means they should not be used interchangeably but that Pentacam and OCT systems were useful additions when looking at preoperative refractive surgery evaluation. When looking at post LASIK eyes (see below, section 2.1.11.1) they found Scheimpflug pachometry maps showed excessive thinning, making OCT a preferable technique for this.

A study comparing 3 different Scheimpflug systems and an OCT system for measuring corneal thickness found they all achieved suitable repeatability for thickness at 5 different sites on the eye [58]. They also found reduced repeatability of measurements further away from the centre of the cornea with all instruments. Slightly better repeatability was found using OCT, the authors attribute this to the

higher scanning speed of the system used and the higher resolution of the system. Their findings were consistent with other studies.

In 2012 a group carried out a review of resolution and repeatability results for different anterior segment biometry measuring devices [59]. They obtained information from the manufacturers of 9 different devices as well as carrying out a literature review. The technologies covered were OCT, Scheimpflug and ultrasound. They found that looking at the manufacturer data the best repeatability for central corneal thickness and minimum corneal thickness was the Galilei system which is a dual channel Scheimpflug camera. This achieved repeatability of $1.38\mu\text{m}$ for vertex thickness and $1.02\mu\text{m}$ for minimum thickness. The next best system was the RTVue OCT system, achieving repeatability of $2.20\mu\text{m}$ and $1.17\mu\text{m}$, and the Artemis ultrasound system, which achieved repeatability of $1.68\mu\text{m}$ and $1.36\mu\text{m}$. Looking at published data they found that there was not always agreement between what manufacturers and authors of studies reported. For Visante, OCT, and SLOCT they found repeatability in published data better than that found by the manufacturer. They conclude this may be due to larger sample size for the manufacturer data. They could find no published data for the Galilei system. No conclusions are drawn by this study but their data indicates that the Galilei system has the highest resolution and repeatability. One disadvantage of that and other Scheimpflug systems compared to OCT systems, is their inability to make certain measurements such as angle to angle distance. This is due to their narrower imaging area. There are other measurements, such as sulcus to sulcus distance, that can only be measured using ultrasound methods since light is unable to penetrate the iris pigment epithelium. No discussion of imaging speed is made in their review.

Overall, anterior segment OCT is a promising technique which has many different uses. Its main strength is the ability to produce high resolution images of the anterior segment without requiring contact with the eye. Other advantages of OCT compared to other systems are good sensitivity and high imaging speed. This means that OCT is a useful technique for using in this study.

2.3 2D Image Segmentation

Image segmentation is the splitting of images into regions that correspond to particular features. In biomedical imaging this mainly involves isolating physiological and biological structures of interest. There are many different techniques available to be used and much work has been done in the field.

The choice of which methods has been motivated by an assessment of their use for segmentation of anterior segment OCT images. Pixel based methods have been investigated for their simplicity and the fact that there have been previous publications segmenting anterior segment OCT images using them. Level set methods using shape have been previously used in retinal OCT images. These have been investigated to see if this work can be adapted to use on anterior segment images. Graph cut methods have been investigated since they have been previously used to achieve a much faster segmentation of images than level set techniques. Finally graph theory segmentation has been investigated since there has been previous work on anterior segment OCT image segmentation using this technique. Each of these techniques will now be discussed in more detail.

2.3.1 Pixel Based Methods

The most straightforward and fastest method of segmentation is based on using individual pixels. A threshold can be applied and pixels are separated into two regions depending on whether they are above or below that value. This is a common method in converting an intensity image to a binary image and is often preceded by other approaches to improve results [60]. Due the fact that segmentation is independent of pixel location this method is sensitive to noise.

It is possible to apply multiple thresholds to an image to slice it into multiple regions if that is required. The key question when using thresholding is how to decide the value of the threshold to use. The speed at which thresholding is carried out makes it possible to manually select a level to use; this however would make the method very time consuming if more than one image had to be segmented. Other simple options include using a fixed value all the time, which would be suitable if all images looked similar, or using the value of a fixed point on an image. Both those methods have the drawback of not being able to deal with varying images.

The most commonly used method for selecting a greyscale value was formulated by Otsu in 1979 [61]. This method selects an optimum value for the threshold based on maximising between class variance. Calculations are made based on the 1D histogram of the image making it a very fast method.

Thresholding struggles to cope with images containing noise. Two regions with a clear difference in mean intensity may have sufficient noise that application of a threshold is unable to segment them effectively. The effect can be reduced by the use of denoising filters but remains a major disadvantage of the method.

2.3.1.1 Improvements to Thresholding

To segment two regions it can be useful to find similarities that can be used to identify the regions. These are sometimes known as continuity based methods since they are trying to segment regions based on similarities within the region. They generally consist of applying a filter to the image and then applying a threshold to segment the image.

The simplest method in this category is to apply a low pass filter to the image. This has the effect of smoothing out noise by applying a Gaussian of a chosen size to each point on the image. For images with a clear difference in mean intensity, but too much noise to be segmented using a simple threshold, this technique works well [62].

Another way of improving segmentation is to use the texture of an image. There are a few different ways to measure this. Small segment Fourier transform, local variance, the Entropy operator, the Laplacian operator, the range operator (difference between maximum and minimum pixels in a neighbourhood), and the Hurst operator (maximum difference as a function of pixel separation) are a few of the most popular types. These allow segmentation of regions with similar mean intensity but different textures. When using this type of method some information is lost on the location of the edges, but since these are often used for images with indistinct edges this is not a major problem. After a texture operator has been applied to the image, it can then be segmented using a threshold as discussed above.

2.3.1.2 Edge Detection

One way of improving pixel based methods is to use edge detection. This method is based on the idea that there are two regions separated by a defined edge. Edge detectors will use either the first or second derivative of the image intensity to segment the image. A simple threshold can be used to only select edges over a certain value, or the method can be incorporated into an active contour model which will be discussed below in section 2.3.2.

Overall, pixel based methods are a very simple method of segmentation. This makes them very fast so they are useful where they are successful, but more robust methods must be found to solve many image segmentation problems.

2.3.1.3 Previous Applications of Thresholding to Corneal OCT Images

Thresholding methods have been used to segment OCT images of the cornea. One group attempting to measure the ocular surface shape SD-OCT images used a

thresholding method [63]. They were interested in looking at the anterior surface over the entire cornea; a custom built SD-OCT machine was used. They found that variations in signal to noise ratio of the images meant that applying a threshold to the intensity image resulted in poor segmentation. The different textures of the features could be utilised to solve this problem, an entropy filter was used to improve results. In their method an entropy map was created, then Otsu's method was used to apply a global threshold to segment the image into two regions. Morphological operations were applied in order to remove small objects and fill in holes. It was assumed that the largest object detected corresponded to the cornea.

Gradient information about the image was then used to improve the segmentation of the image. The gradient of the image was calculated and the point with maximum gradient was assumed to be the boundary of the anterior surface. The brightest pixel in regions along the initial boundary was selected and used to specify a new edge. A 10th order polynomial was then fitted to this to smooth the identified boundary. Corrections were also applied to take into account distortions due to the OCT method. The study managed to achieve good repeatability of their results when comparing segmentation of several different images of the same eye. However the study did not compare their segmentation technique to expert manual segmentation, so it is difficult to assess the accuracy of segmentation. The study was also limited to identifying the anterior boundary of the cornea only.

Another study attempted to use thresholding techniques to calculate the anterior chamber angle [43]. Here the important structure to identify is Schwalbe's line which forms part of the boundary of the cornea and the anterior chamber. A threshold was used to segment the image and morphological operations used to remove speckle noise from the image. They assumed the largest object found was the cornea and checked to see if the second largest object was big enough to be the iris. A 4th order polynomial was fitted to the line on the posterior boundary of the cornea. Schwalbe's line is characterised as being a disruption to this fit. Dewarping was performed to remove distortion due to the OCT scanning method. Comparison of their results to segmentation by experts showed their results worked just as well as manual labelling in 37 out of 40 images. The failure in some images was due to heavy noise around the angle recess. A Zeiss Cirrus HD-OCT system was used to acquire images. This was able to achieve a transverse resolution of 15 $\mu\text{m}/\text{pixel}$ and axial resolution of 5 $\mu\text{m}/\text{pixel}$. A wavelength of 840 nm was used, the system

had a shorter penetration depth than the Visante OCT system, meaning it is only capable of imaging part of the cornea.

Overall, thresholding has been shown to have some promising signs. There are however a number of problems with it, the main one being its heavy sensitivity to noise. So far studies using it on the human cornea have only been done with small amounts of data. If large scale trials were made they would involve more images with lower signal to noise ratio (SNR) but this would cause thresholding to break down. The quick and simple nature of the method makes it useful as a starting point for segmentation but a more complex method may be needed to produce consistent reliable results.

2.3.1.4 Central Corneal Thickness

There have been several of papers published of studies looking at central corneal thickness using OCT. Many studies measuring anterior segment dimensions have used manual annotation [42, 64]. Methods for automatic segmentation have also been developed [31]. This is done by examining the intensity profile of a cross section of the image. Locations a fixed distance on either side of the central noise artefact are used. The first intensity peak corresponds to the location of the front surface of the cornea, the second to the location of the epithelium-Bowman layer and the last peak to the endothelium – aqueous interface. The distances between these peaks are then used to measure central corneal thickness and epithelium thickness. This method has been shown to work on a variety of different OCT devices [29] with results agreeing with those achieved through manual segmentation. A median filter is used to pre-process images for this technique

This technique has been shown to work well when looking at central corneal thickness but how far it can be extended to cover the entire cornea is unclear. The centre of the cornea is usually the location of the best image quality.

2.3.2 Active Contours

Since their introduction in 1988 [65], active contour based methods have been extensively used in image segmentation. There are many different formulations of this technique, depending on how it is wished to separate the regions. Active contours work by having deformable contours that segment the image into regions. This can be driven by regional based statistics, edge information, or using predefined shape information. Various regularising forces can also be added to keep contours smooth, push them past noisy edges, and overcome other problems that may be encountered.

2.3.2.1 Level Set

Level set based formulations are a widely used technique in the field of image segmentation. These are where the contour is represented as the zero level set of a smooth function, often known as a level set function [66]. The evolution of this surface is represented as the minimisation of an energy function with terms added to produce the desired results. One of the advantages of the level set formulation is that topological changes such as breaking or merging sections are well defined and will occur as and when required.

The boundary between two regions is defined by $\Phi = 0$; inside the region $\Phi > 0$ and outside $\Phi < 0$. The motion can be analysed using a velocity field v where $\frac{d\Phi}{dt} + v\nabla\Phi = 0$. The velocity field is determined by the energy function and governs how the level set will evolve with subsequent iterations [67].

2.3.2.2 Edge Detection

Edge detection is based upon detecting an edge using the gradient of the image. It can be incorporated into level set methods. For example the Caselles method [68] uses the following energy equation

$$E(T) = \int_0^1 g(|\Phi(q)|)|T(q)|dq \quad (7)$$

where $\Phi(q)$ is the level set function, and

$$g(T) = \frac{1}{1 + |\nabla(G * I)|^2} \quad (8)$$

where I is image intensity, T the curve, and G a Gaussian filter. The image has been convoluted with a Gaussian to smooth it, then the gradient of image taken. The evolution equation for this is

$$\frac{\partial\Phi(x)}{\partial t} = g(x)|\nabla\Phi(x)|(c + \kappa) + \nabla g(x)\nabla\Phi(x) \quad (9)$$

where κ is the curvature of the evolving contour, x is the position of a point in the image, and c is a constant that acts as a balloon force. Increased gradient will slow movement of the curve, causing the contour to stop at regions of high gradient. The balloon force can be adjusted to affect the sensitivity of the edge detection. This method suffers when there are noisy images since the contour may get stuck on regions of high gradient due to noise.

These methods have been used for studying OCT images of the retina but extensive pre-processing is required to reduce the effect of speckle noise [69].

2.3.2.3 **Active Contours without Edges**

Using image gradient information has problems segmenting noisy images or those without sharp edges. An alternative technique is to use regional statistics to segment an image instead. This method was first proposed by Chan and Vese in 2001 [70]. In the simplest case segmentation between two regions is carried out; the two regions are assumed to have approximately constant intensities within the region and a substantial difference to those outside the region. The energy term then consists of two terms F_1 and F_2 with

$$F_1(c) + F_2(c) = \int_{inside(c)} (u_0(x, y) - c_1)^2 dx dy \tag{10}$$

$$+ \int_{outside(c)} (u_0(x, y) - c_2)^2 dx dy + \int_{\Omega} \delta(\Phi_i(x, y)) |\nabla \Phi_i(x, y)| dx dy$$

where C is a variable curve, u_0 is image intensity, and constants c_1 and c_2 are the averages of u_0 inside C and outside C respectively. This energy function is then minimised; this will equal zero when C divides two regions of equal intensity. Since images requiring to be segmented are nearly always more complicated than two uniform layers the addition of regularisation terms is needed. The third term is a smoothing term to keep the contour a smooth curve; a function of length of the contour is used for this purpose.

2.3.2.4 **Shape Constraints**

The segmentation of images can be improved by using information about an object's shape. Often the general shape of the object to be segmented is known and incorporating this information into the energy function can improve results [71]. There are two main ways to incorporate shape information. One is to define a shape an object should look like, for example a circle or an ellipse. The other is using statistical information.

2.3.2.4.1 Statistical shape models

Statistical models require a set of training images [72]. These can then be used to add another energy term to the function. This must normally be done in addition to a regional based term based on the Chan Vese formulation discussed above [73]. The energy formulation for this method can be best seen through a Bayesian

formulation of the problem [74]. The level set function is denoted ϕ , image information I , with training shapes being $\tilde{\Phi}$ and \tilde{I} respectively [75]

$$\tilde{\Phi} = \arg \max p(\Phi | I, \tilde{\Phi}, \tilde{I}) \quad (11)$$

$$\tilde{\Phi} = \operatorname{argmin}(-\log(p(I|\Phi, \tilde{\Phi}, \tilde{I})) + \log(p(\Phi|\tilde{\Phi}))) \quad (12)$$

$$\tilde{\Phi} = \operatorname{argmin}(E_{image} + E_{shape}) \quad (13)$$

The image information I can be split into three regions; h_{in} inside the contour, h_{out} outside the contour, and h_{en} over the entire region. h_{in} and h_{out} both depend on ϕ whilst h_{en} does not. There is a simple relationship between the three

$$h_{en} = \alpha h_{in} + (1 - \alpha)h_{out} \quad (14)$$

where α is the ratio of the area inside evolving contour to entire area. These expressions can be used to split the image energy term to give

$$\begin{aligned} \tilde{\Phi} = \operatorname{argmin} &(-\log(p(\alpha(\Phi)|\tilde{\alpha})) - \log(p(h_{in}(\Phi)|\tilde{h}_{in})) \\ &- \log(p(h_{out}(\Phi)|\tilde{h}_{out})) - \log(p(\Phi|\tilde{\Phi}))) \end{aligned} \quad (15)$$

$$= \operatorname{argmin}(E_{\alpha} + E_{h_{in}} + E_{h_{out}} + E_{shape}) \quad (16)$$

These expressions are substituted into the previous equation which is minimised to find the best level set to segment the image.

This method of using a statistical shape force to help segmentation allows segmentation of images where the object to be segmented is not well differentiated from the surroundings by either edge information or by regional intensity variation. This method has been used successfully in a variety of medical imaging problems and has been shown to perform better than methods not using shape energy [76-78].

2.3.2.4.2 Parametric Shape model

An alternative way to incorporate shape energy is to use a predefined shape instead of training data. This has the advantage of not requiring the inputting of large amounts of training data to the system. One area where this has been used successfully was in segmentation of different layers of the retina in OCT images of rats eyes [79]. The shape term is needed in this case due to difficulties caused by speckle noise inherent to OCT and blood vessel artefacts within the retina.

The different layers of the retina were modelled as a series of concentric circles, since this roughly corresponds to the retinal structure. In this method a level set formulation was used. In order to avoid the need for reinitialisation of the level set a regularising term was added to ensure the level set remained a valid level set. This was based on work by Li et al. [80]. A contour length term was also used making the regularisation term

$$E_R = \sum_{i=1}^{R-1} \int_{\Omega} \delta(\Phi_i(x, y)) |\nabla \Phi_i(x, y)| dx dy + \sum_{i=1}^{R-1} \int_{\Omega} \frac{1}{2} (|\nabla \Phi_i(x, y)| - 1)^2 dx dy \quad (17)$$

where $\Phi_i(x, y)$ is the level set function and $\delta(\Phi_i(x, y))$ is a delta function that is 1 on the boundary and zero elsewhere. The first term here is the contour length term and the second the level set regularisation term.

The shape energy term was based on the distance to the circle from a point defined as

$$D_i(x, y) = ((x - c_x)^2 + (y - c_y)^2 - r_i^2)^2 \quad (18)$$

where (c_x, c_y) is the centre point common to all the circles and r_i is the radius of circular prior the i th interface. The energy term to encourage the level set to lie on the circle is then

$$E_S = \sum_{i=1}^{R-1} \int_{\Omega} D_i(x, y) \delta \Phi_i(x, y) |\nabla \Phi_i(x, y)| dx dy \quad (19)$$

The $\delta \Phi_i(x, y) |\nabla \Phi_i(x, y)|$ term ensures that the shape term only has nonzero values on the zero level set. E_S is minimised when $\Phi_i(x, y)$ lies exactly on the circular shape. The shape parameters are estimated using a least squares fit, with ϕ held fixed.

In addition to the shape and regularisation terms they also used regional intensity terms from Chan and Vese method same as equation 10 here. Their energy function was a combination of equations 10, 27 and 19 giving

$$E = \lambda_L E_L + \lambda_S E_S + \lambda_R E_R \quad (20)$$

where E_L is energy term from image intensity (given in equation 10), E_S is shape energy term from equation 19 and E_R is regularisation term from equation 17. Each

term has a corresponding λ value which is a coefficient which determines the relative strength of the different terms.

The optimisation of the relative weighting of the different energy terms is important in achieving good results. This study used adaptive weights, starting off with the regional Chan Vese term dominating and then applying a shape term later on in segmentation. They were able to successfully segment several different layers in the retina. Results achieved were better than those using just the Chan Vese active contours without edges method, and those using the shape term but not adaptive weighting of the different terms.

A similar method was used to segment cardiac magnetic resonance (MR) images [81]. An elliptical shape term was used for this method; in this case both regional statistics (active contours without edges) and edge based terms were used. Similar adaptive weighting of relative terms was used here; initially the region and edge based terms were allowed to dominate, then the shape term was allowed to dominate in the final iterations.

2.3.2.5 *Relevance to this Work*

Level set techniques with shape have been previously used to segment retinal OCT images successfully. For this application, the main problem is regions of low SNR within the cornea. This means that using a technique that includes information on the expected shape of the cornea could be useful.

2.3.3 Graph Cut Segmentation

Level set methods have many advantages; they are easily able to cope with changes in topography of segmentation, they are global methods with minimal dependence on starting conditions and have been shown to produce good results. There are however some limitations to these techniques. The primary one is speed. Graph cuts have been shown previously to be a quick alternative to level sets [82] in image processing.

Graph cuts are a way of segmenting an image exploiting a min-cut algorithm for graph labelling [83]. A global optimum cut can be obtained for an image. A set of pixels with a corresponding set of labels are input and the goal is to minimise an energy function. The motivation for using these techniques is that there are very efficient max flow algorithms that can be used to minimise the energy [84]. These methods have been used increasingly in image segmentation and other image

processing problems. In order to use graph cuts an energy function must be constructed in the correct form [85].

The most basic formulation of the problem that is faced is to solve a pixel labelling problem. If we have a set of pixels P with a corresponding set of labels L . A labelling x can be found which minimised some energy function. The standard form of the energy function is

$$E(x) = \sum_i E_i(x_i) + \sum_{i,j \in N} E_{ij}(x_i, x_j) \quad (21)$$

where n is a neighbourhood system on pixels, $E_i(x_i)$ is a function that measures the cost of assigning label to pixel x_i to pixel i , and $E_{ij}(x_i, x_j)$ is a function that measures the cost of assigning labels (x_i, x_j) to adjacent pixels i, j [86]. These are given by

$$\begin{aligned} E_i(x_i) &= E_i^0(1 - x_i) + E_i^1 x_i \\ E_{ij}(x_i, x_j) &= w_{i,j} \left((1 - x_i)x_j + x_i(1 - x_j) \right) \end{aligned} \quad (22)$$

This energy function can be minimised by graph cuts if and only if $w_{i,j} \geq 0$ for any $(i, j) \in N$.

Implementation of Chan and Vese active contour without edges has been done using graph cuts several times previously showing ~1,000 times faster results [87] when compared to level set implementation. For this model the region is $E_i(x_i)$ since it is a global function. The curvature term is then represented by $E_{ij}(x_i, x_j)$. There are different ways to represent the curve length in graph cut optimisations. A commonly used method is the Cauchy- Crofton formula.

2.3.3.1 **Cauchy-Crofton Formula**

The graph cut will divide the pixels into two different sets, an image set and a background set. For a given cut the energy cost can be defined as

$$|C|_G = \sum_{(i,j) \cap C \neq \emptyset} w_{i,j} \quad (23)$$

Boykov and Kolmogorov [84] introduced the cut metric and related length through the Cauchy Crofton formula. If a set of lines in a plane is defined by (ρ, θ) where $x \cos(\theta) + y \sin(\theta) = \rho$, the length of a curve C is given by

$$length(C) = \frac{1}{2} \int_0^\pi \int_{-\infty}^\infty n_c(\theta, \rho) d\rho d\theta \quad (24)$$

where $n_c(\theta, \rho)$ is the number of interactions between the line $x\cos(\theta) + y\sin(\theta) = \rho$ and the curve C. This gives

$$length(C) = \sum_{k=1}^n n_c(k) \frac{\delta^2 \Delta\theta_k}{2|e_k|} \quad (25)$$

Different neighbourhood schemes can be used to approximate the length. Increased size of neighbour generally produces to a better result but will take longer to carrying out, n_c is given by neighbourhood scheme. δ is the grid size, $\Delta\theta_k$ the angular difference between nearest grid lines and $|e_k|$ is the length of the edge e_k . Edge weights can be chosen to be

$$\omega_k = \frac{\delta^2 \Delta\theta_k}{2|e_k|} \quad (26)$$

This can be substituted back into equation 21 to give.

$$|C|_E = \sum_{k=1}^n n_c(k) \omega_k = \sum_{i=1}^{N_E} \omega_i \quad (27)$$

This means $|C|_E$ can be used as the cost of $|C|_G$ in equation 21 of a cut in graph G.

2.3.3.2 **Shape in Graph Cut Segmentation**

There has been previous work adding shape terms into graph cuts models. One study used an elliptical shape prior to improve segmentation of lymph nodes in pelvic MR images and for detection of human faces in images [88]. They used an image based energy term from Chan and Vese with

$$D_p(object) = |I(p) - u_i| \quad (28)$$

$$D_p(background) = |I(p) - u_o| \quad (29)$$

where $I(p)$ is the image intensity at pixel p , u_i is the mean intensity inside the curve and u_o the mean intensity outside the curve.

An additional shape prior term, based on a binary shape mask M valued at 0 inside an ellipse and 1 outside, is used. The shape based term is

$$D_p(\text{object}) = |M(p) - 1| \quad (30)$$

$$D_p(\text{background}) = |M(p) - 0| \quad (31)$$

These were taken together to equal $E_i(x_i)$ in equation 21. The second term in 21 is not changed by addition of shape and so a similar term to that used in Danek's paper is used [89]. This gives an energy function of

$$E(x) = \sum_i |I(p) - u_i| + |M(p) - 1| + \sum_{i,j \in N} |C|_E \text{ inside}, \quad (32)$$

$$E(x) = \sum_i |I(p) - u_0| + |M(p) - 0| + \sum_{i,j \in N} |C|_E \text{ outside}.$$

2.3.3.3 Using a Shape Variability Model

Ali et al. [90] developed a method for segmenting images of a kidney which exploited shape information. They used a set of training data to produce a distance probabilistic model. A method was produced that combined this shape information with image information to produce a graph cut framework. Their technique was shown to work much better than those without the shape term.

The energy function they used is defined by

$$E(\mathbf{f}) = \sum_{p \in P} S(f_p) + \sum_{p \in P} D(f_p) + \sum_{(p,q) \in N} V(f_p, f_q) \quad (33)$$

where $S(f_p)$ is the shape term explained in the next section, $D(f_p)$ is the image information similar to the one used previously, and $V(f_p, f_q)$ is the smoothness term.

The shape model they used is based on using a set of training data. The first step is to align the segmented images using 2D rigid registration. These images are then converted to binary images. A three stage shape is then created, $P_s = O \cup B \cup V$ where O is the object, B is the background, and V is the variability region. A distance probabilistic model is used to describe the variability region. The probability of a pixel to be in the object decreases exponentially as the discrete distance increases. The distance histogram can be modelled as a Poisson distribution. The histogram entity at distance d_p is given by

$$h_{dp} = \sum_{i=1}^M \sum_{p \in C_{dp}} \delta(p \in K_i) \quad (34)$$

where delta function $\delta(p \in K_i)$ is 1 when $(p \in K_i)$ is true and zero otherwise, M is the number of training images, and K_i the kidney region in the training image i . The distance is changed till the whole distance domain is covered. This is then multiplied by the kidney prior value defined as follows

$$\pi_K = \frac{1}{M|V|} \sum_{i=1}^M \sum_{p \in C_{dp}} \delta(p \in K_i) \quad (35)$$

The shape energy term is defined as $S(f_p) = -\ln P(d_p|f_p)$ where the distance marginal density of each class $P(d_p|f_p)$ is shown below.

There is a slight deviation between the estimated and empirical densities since f_p does not follow a perfect Poisson distribution. The deviation is modelled as a linear combination of discrete Gaussians with +ve and -ve components. This gives the following

$$P(d_p|f_p) = v(d_p|\lambda_{f_p}) + \sum_{r=1}^{K_f^+} w_{f_p,r}^+ \gamma(d_p|\theta_{f_p,r}^+) - \sum_{l=1}^{K_f^-} w_{f_p,l}^- \gamma(d_p|\theta_{f_p,l}^-) \quad (36)$$

where $v(d_p|\lambda_{f_p})$ is a Poisson density with rate λ , and $\gamma(d_p|\theta)$ is a Gaussian density with parameter $\theta \equiv (\mu, \sigma^2)$ with mean μ and variance σ^2 .

Ali et al. reported that this model showed improved results compared to segmentation techniques not including shape information.

2.3.3.4 **Relevance of Graph Cut Segmentation to This Work**

Graph cut techniques have been shown by other groups to be capable of achieving much faster segmentation than level set methods. For the segmentation of anterior segment OCT images speed is an advantage for two reasons. Rapid segmentation will allow for real time segmentation, which means that patients can have imaging carried out and then discuss the results of the imaging with a doctor at the same visit. The other reason is the development of 3D imaging. 3D imaging will give larger amounts of data so rapid processing is needed in order to make this practical.

2.3.4 Graph Theory Based Methods

Another method of segmentation of images is through the use of graph theory based techniques. These work by setting weights for joining different pixels to each other and then finding the lowest weighted path through an image [91]. Typically edge information is used to determine the weighting of the path.

In general the image is represented as a graph of nodes, with each node corresponding to a pixel. These nodes are connected by edges, and a connected set of edges form a pathway. The edges can be assigned different weights to create a preferred route. With appropriate weights assigned, Dijkstra's algorithm [92], a simple and fast technique, can be used to determine the lowest weighted path. Graphs can be set up when either the edges have weightings or the nodes contain the weighting information.

An alternative method to find the shortest path through the graph is to use dynamic programming. The principle behind this method is the idea, that whatever the path from start point A to node D, there exists an optimal path between D and the endpoint. This means if the optimum start point to end point goes through D, then both sections start point to D and D to end point are also optimal paths.

An example graph segmented using this technique is shown in Figure 9 below.

Figure 10 explains how the technique works on a step by step basis.




Figure 9: Example of a graph searching sequence. Graph to be segmented, each point on the graph can be taken to represent a pixel in an image. The value of each node on the graph is determined by the cost function.




Figure 10: Step by step showing segmentation of graph using dynamic programming. The first step is to find the shortest path to the second layer of the graph from an arbitrary starting point. Cumulative cost is shown in small numbers above each node. The second step is to do this again for the third layer. Cumulative cost to get to each node is shown on the diagram. The third step does this again for the fourth and final layer. The node with the shortest path is marked. The fourth and final step is to trace back from the shortest path node to find the optimal path back to the start.

In the example shown above each point was allowed to connect three possible points in a forward direction. The connectivity of the graph can be altered depending on the problem that is being solved. The cost for each node in the above graph is given by the expression

$$C(x_i^{m+1}) = \min_i (C(x_i^m) + g^m(i, k)) \quad (37)$$

where $C(x_i^{m+1})$ is the cost associated with the node x_i^{m+1} , and $g^m(i, k)$ is the cost to get from node x_i^{m+1} to node x_i^m . For the entire problem this becomes

$$\min(C(x^1, x^2, \dots, x^M)) = \min_{k=1, \dots, n} (C(x_k^M)) \quad (38)$$

where x_k^M are the end nodes, M is the number of layers between start and end points, and $C(x^1, x^2, \dots, x^M)$ is the cost of a path between the first and last graph layer. The graph must be correctly constructed to allow for this optimisation to take place.

In order to segment images using this method it must first be determined how to set the weightings for different paths. Examples of general techniques that can be used for this are discussed by Sonka et al in their book [93]. They discuss four principle methods.

- 1) Strength of edges of the boundary. This is the idea that the stronger the edges of the boundary, the more likely this is to be the true boundary.
- 2) Curvature of boundary. In some applications boundaries with small curvatures are preferred. The total curvature can be evaluated as the sum of local curvature increments

$$diff(\Phi(x_i) - \Phi(x_j)) \quad (39)$$

where *diff* is a function evaluating the difference in edge direction between two consecutive border elements.

- 3) Proximity to approximate boundary location. If some information is known about where the boundary should be then paths which are closer to this can be favoured. An element can be weighted based on its distance from the estimate boundary location.
- 4) Estimates of the distances to the end point. If a reasonably straight boundary is wished for, then elements closer to the end can be favoured.

These techniques have been used for a wide variety of purposes and the function used to determine the cost should be varied according to the desired solution. Often good cost functions will combine several different components of these techniques to give a solution.

2.3.4.1 **Graph Theory Segmentation of Retinal and Corneal OCT images**

Graph theory techniques have been used to segment retinal and corneal OCT images [94-97]. Vertical gradient was used to set the cost function. One group carried out this technique using SD OCT images of the retina [94]. They first flattened the image to improve results. In order to flatten the image they assumed that the Retinal Pigment Epithelium (RPE) layer was the brightest pixel in each column after denoising with a Gaussian filter. The image was then adjusted to make this layer flat in the image. This was done to reduce errors due to the segmented boundaries taking shortcuts to the edge of the image. Weights were calculated using vertical intensity gradients

$$w_{ab} = 2 - (g_a + g_b) + w_{min} \quad (40)$$

where w_{ab} is the weight of the edge connecting a to b , g_a is the vertical gradient at node a , g_b is the vertical gradient at node b , and w_{min} is a small minimum value added for stability. The gradient must first be normalised to be between 0 and 1 before this can be used. The direction of the gradient can also be taken into account and used to identify particular boundary line in the image. The group used Dijkstra's method [92] to find the shortest path through their image. This normally requires the setting of start and end points for the segmentation. To avoid having to do this in their application, they added a zero weight column to each edge of the image so the segmentation was able to start at any row on the first column of actual image data.

Their method was able to segment 7 different boundaries in the retina; in comparison with expert manual segmentation the method performed well, having a lower mean difference when compared to manual segmentation than when comparing two expert graders.

Another group tried a similar method also looking at retinal OCT images [95]. This group was aiming to segment 9 different layers of the macular. Part of their focus was on improving speed compared to previous studies, since they believed the slow speed of previous methods made them unsuitable for clinical use. Their method used a combination of Canny edge detectors and dynamic programming to segment nine layers. The weighting function they used was given by

$$C(i, j) = w_1 * Canny(i, j) + w_2 * Axial(i, j) + w_3 * Others(i, j) \quad (41)$$

where *Canny* is a modified Canny edge detector which evaluates the strength of different edges, *Axial* is an axial intensity gradient map which uses a larger filter size to gain additional information and *Others* varies for the different boundaries and is not specified in their paper, other than saying it may include intensity information. $w_{1,2,3}$ are weighting terms which govern which of the three terms are most important. In order to detect the different boundaries, they varied the threshold for their Canny edge detector depending on the boundary they were trying to find.

They used a different method to find the shortest path once they had set the node weight values, choosing a method based on dynamic programming. This has been used extensively on different image segmentation problems including others looking at retinal OCT images [98]. Here the minimum cost to reach position (i, j) is defined by

$$t(i, j) = \begin{cases} \infty & j < 1, j > m \\ C(i, j) & i = n \\ \min_{p=j-2:j+2} (t(i-1, p) + C(i, j)) & \text{otherwise} \end{cases} \quad (42)$$

where $t(i, j)$ is the minimum cost to reach point (i, j) , $C(i, j)$ is the node cost function, i is the X direction index with maximum n , and j the Y direction index with maximum number m . This method is an efficient way of finding an optimum path where the start and end points are unknown.

The first group also used the same graph based method to look at OCT images of the cornea [97]. The images were acquired using a SD OCT system manufactured by BiopTigen. They were attempting to segment three boundaries in the cornea; the air-epithelium interface, the endothelium-aqueous interface, and the Epithelium-Bowman's layer interface.

The first step of their method was to remove several artefacts from the image. Central saturation artefacts at the corneal apex and lower SNR away from the centre are caused by the telecentric scan pattern used by most systems. Horizontal line defects are also often present, which is due to the effect central saturation has on the DC subtraction algorithm used by many SD OCT systems. The horizontal artefact can be removed by subtracting the mean intensity of each row from every pixel in that row. This works best if the horizontal artefact is of uniform intensity and does not pass through the cornea. The vertical artefact is detected by first applying a vertical median filter to accentuate it, then looking for abrupt changes to the mean intensity of A scans. Once the central artefact is detected its effect on segmentation can be removed.

The image gradient was used to produce path weights between points on the image. The weighting function used was identical to that they had previously used on the retina see equation 40.

Dijkstra's method [92] was implemented to find the lowest weighted path across the image. Problems were encountered in low SNR regions where the tendency for Dijkstra's method to find straight lines overcame the weighting of the correct path on the image. It was assumed that the cornea boundary corresponds to a 2nd order polynomial and the curve from the central region was expanded into regions of low SNR where the method had failed. The segmentation is then adjusted as it was segmented based on maximum gradient; however the actual boundary is lower than this at the point of maximum intensity. Therefore the segmentation is adjusted to find the lowest weight path, this time using intensity and not gradient to assign

lengths, within 20 μm from the initial path. The lower boundary was found by assuming the corneal thickness at the apex is approximately the minimum thickness across the entire cornea. The image was flattened based on the detection of the upper boundary and the same method used, limiting the area path could travel through by the previous assumption. This study managed to achieve results comparable to manual grading with 40 B scans. The mean difference between automatic and manual segmentation of the anterior boundary was 0.6 pixels compared to 1.5 pixels between two images manually segmented by different experts.

This method has been shown to work for healthy corneas but has difficulties detecting boundaries in lower SNR regions. The group indicated they were trying to extend their work to cover a wider area, but the method described above uses a polynomial to extend the segmentation beyond the central region, assuming the cornea is regular. It was based on using SD OCT images from a commercially available system; it was not capable of taking an image of the entire anterior segment.

2.3.5 Relevance to this Work

Image processing is a wide and varied field with many different techniques available. Many of these have been used previously to segment retinal OCT images. There has been less work looking at the segmentation of anterior segment OCT images. Most groups that have previously looked at the cornea have only focussed on part of it rather than the entire anterior segment. One group managed to use pixel based methods to segment the entire ocular surface. This group used a custom built SD OCT machine, but their method seems to be a good starting point for developing a more robust method. Methods involving use of level sets, graph cut and graph theory based techniques have also produced promising results previously.

2.4 3D Segmentation

Carrying out imaging of structures using only 2D images is not sufficient for all applications. A 2D scan only covers results from a particular view of a 3D object. For some imaging modalities this means that objects may be hidden behind other objects, other imaging modalities will only acquire a 2D single slice of the object. So while this will show detail on that particular area it does not give an accurate showing of the entire object. OCT imaging has generally been developed in a way

to give 2D images of objects. It is possible to extend this imaging system to 3D by acquiring multiple 2D images and combining them to give a 3D model. This then poses two challenges for 3D image segmentation; the first is an alignment problem and the other is the actual segmentation itself.

The problem of image alignment is due to the fact what is being imaged is live biological tissue. This makes it susceptible to movement artefacts which in the case of the eye are often involuntary and so it is not possible to get the patient to remove these. In order to reduce the effect of movement artefacts either an increased imaging speed or image registration techniques can be used.

2.4.1 3D Retinal OCT Imaging

OCT imaging is used extensively to study the retina. Initially this was done using 2D scans but the development of faster spectral domain OCT imaging led to extensive take up of 3D OCT imaging of the retina. Looking at the registration and segmentation techniques used will be helpful for investigating the cornea.

Generally a raster scan pattern is used on the retina. This is where a series of parallel scans at a known distance apart are used to image the region of interest. One of the problems faced in ophthalmic imaging is involuntary eye motions [99]. These are present even in high speed SD OCT images and have to be accounted for in the image alignment process.

This can be done by generating two images simultaneously. A fundus image of the retina is acquired at the same time as the 3D raster scan. A fundus image is a 2D face-on image of the retina. This can be used to visualise the location of key features in the retina such as blood vessels. Blood vessels can be detected in both the fundus image and OCT image. These can be used as landmarks to allow for alignment of the images [100].

This method has been adopted for commercially available retinal OCT techniques. However it has limited applicability for OCT images of the cornea. The cornea lacks any such obvious features that pass through the image, so this method cannot be used.

When carrying out 3D segmentation two methods can be used, either true 3D segmentation where the entire volume is considered at the same time, or “2.5D” segmentation where a series of 2D scans are segmented separately before being

combined to produce a 3D image. Which is chosen depends on the method of segmentation that is desired to be used.

Level set methods of segmentation can be expanded into 3D with trivial changes to the mathematical formulation. The increasing number of data points when going to 3D volumes does result in practical limits in terms of computational power and speed of the segmentation technique. For example one group reported using a 3D level set based method with shape prior to segment 172x148x124 voxel MR brain images [101]. They gave a stated time of “less than an hour” for the 3D images compared to 2 minutes for 2D slices of the same data.

Graph theory segmentation techniques cannot easily be formulated into a natural 3D method. This means work with this method can be done more easily using a 2.5D segmentation method. For example, a group using graph theory to segment nine layers in the retina [95] segmented 2D images separately. They managed an acquisition time of 16 seconds for a 480x512x128 voxel volume.

Graph cut methods can be expanded to be fully 3D; they just require the construction of a suitable graph. One group used a graph cut method to segment 5 different layers in the retina [102]. They were using images acquired from 6 radial scans of the retina. A multistep process was used in their segmentation. First they aligned each image so that the RPE (an easily definable boundary) was horizontal. They then performed a 2D segmentation step on all the images, finding 3 different boundaries. Using this segmentation a registration step was performed using mutual information metric to align all the images to the first image taken. This then gave them a series of 2D scans aligned in such a way that they could carry out 3D segmentation on them. They then found the optimum segmentation by finding the min cut of a 3D geometric graph they constructed.

They defined a surface as a function $f(r, \theta)$ mapping (r, θ) pairs to z values. They applied smoothness constraints to the surfaces, only including feasible surfaces where the difference in z values of neighbouring points is less than a certain threshold. They then constructed a 3D geometric graph for each surface; costs were assigned to each vertex and then the minimum-cut closed set was found. The cost function used varied, depending on the different surface they were trying to detect. They successfully managed to detect surfaces just as accurately as two different human observers could detect them. No information was given on how long it took for their programme to run.

This method could be applicable to our work, since they have used a similar scan pattern to the one we have used. They have also used a very simple registration technique that relies on accurate detection of boundaries and uses them to align images.

2.5 Image Alignment Methods

2.5.1 Iterative Closest Point

Iterative closest point algorithm (ICP) is a fairly simple image registration method. This automatically registers a set of points from one surface to another [103]. It assumes that the surface being registered is a sub set of the other, so only one surface can contain points that do not exist on the other surface. There is a likelihood of the solution getting stuck in a local minimum, so a good initial estimate of the solution is needed [104].

ICP has been used extensively in different image registration techniques. It works by carrying out a number of steps in an iterative fashion in order to improve the registration of two data sets [105]. The first step is to associate all the points within the object to a point in the reference template. This step is carried out many times so it is important that it is done in an efficient way. One method uses a linear technique where an exhaustive search of all possible connections is carried out. This works best for problems with a small number of points or a high number of dimensions. An alternative way of finding the nearest neighbour is to use a space partitioning search. This reduces the space where the points are searched for. This is more complex in relation to dimension rather than number of points, so is faster for problems with a large number of points and a low number of dimensions.

Once the points are all paired up, the distance between the different sets of points is calculated. A transformation that reduces this distance is then found. The process is then repeated in an iterative fashion until an optimum transformation is found. The stopping criteria usually used is either to have a maximum number of iterations that cannot be exceeded or to stop when the distance between the points and the model falls below a certain threshold value.

2.5.2 Genetic Algorithms

Finding the optimum value for a function is a general problem that is often encountered in image processing problems. The problem in a general form is; for a finite domain D a function f is defined such that $f:D = R$, R being the set of real

numbers, to find the best value of D under function f . This best value is the value $x \in D$ giving either a minimum or a maximum of function f :

$$f_{min}(x) = \min_{x \in D} f(x), \quad f_{max}(x) = \max_{x \in D} f(x). \quad (43)$$

The function f is known as the objective function. Genetic Algorithms (GA) are a method for finding the minimum (or maximum) of this function. In order to achieve success in the task the objective function must well represent what is wished to be found. Therefore care must be taken in the construction of the objective function.

The simplest approaches to optimisation of the function use calculus based techniques such as gradient descent. The gradient of the objective function gives the direction of steepest descent and this can be followed to find a minimum. The main problem with this method is that a local minimum can easily be found and there is no easy way to find a global minimum. An example of a function with several local minima can be seen in Figure 11 below. Genetic Algorithms represent an alternative technique which attempts to find a global minimum by introducing a random element to the search.




Figure 11: Example of a function containing several local minima. Limitations of gradient descent could cause solutions to be stuck at $x = 8$ or at $x = 5.2$ where local minima are found

2.5.2.1 *How Genetic Algorithms Work*

The central idea behind genetic algorithms (GA) is using natural evolution mechanisms to search for the minimum. The use of GA does not guarantee that the global optimisation is found but empirical results show that the final solution is normally very close.

GA uses a population of possible solutions. The value of the objective function at each of these points is calculated; the value of this function is known as the fitness function in GA. Each parameter can be represented as a binary string. It is these strings rather than the parameters themselves that the method processes to find the solution. It is also possible to use real valued GA, which is where the parameters themselves are used directly in the GA rather than binary strings. This has the advantage of not requiring switching between binary strings and real values when calculating the cost function.

In order to generate a new population from the previous one, GA uses three processes known as reproduction, crossover and mutation. The idea behind this is that a solution with a better value for the fitness function will have a better chance of surviving.

The reproduction step is based on the idea of survival of the fittest using a probabilistic treatment. First all members of the population are ranked according to their value of the fitness function. Often a number of elite children are copied through without change, ie. the best 3 or 4 individuals from the population are included in the new population so as not to lose the best fit that has been found so far. The rest of the population is made up of children generated by the processes of crossover and mutation. Strings with a better fitness value have a higher chance of being included in the new population. This can result in there being multiple copies of the same string in the new population if they have a good fitness value. Figure 12 below shows the three different methods by which reproduction occurs; crossover and mutation are explained below.

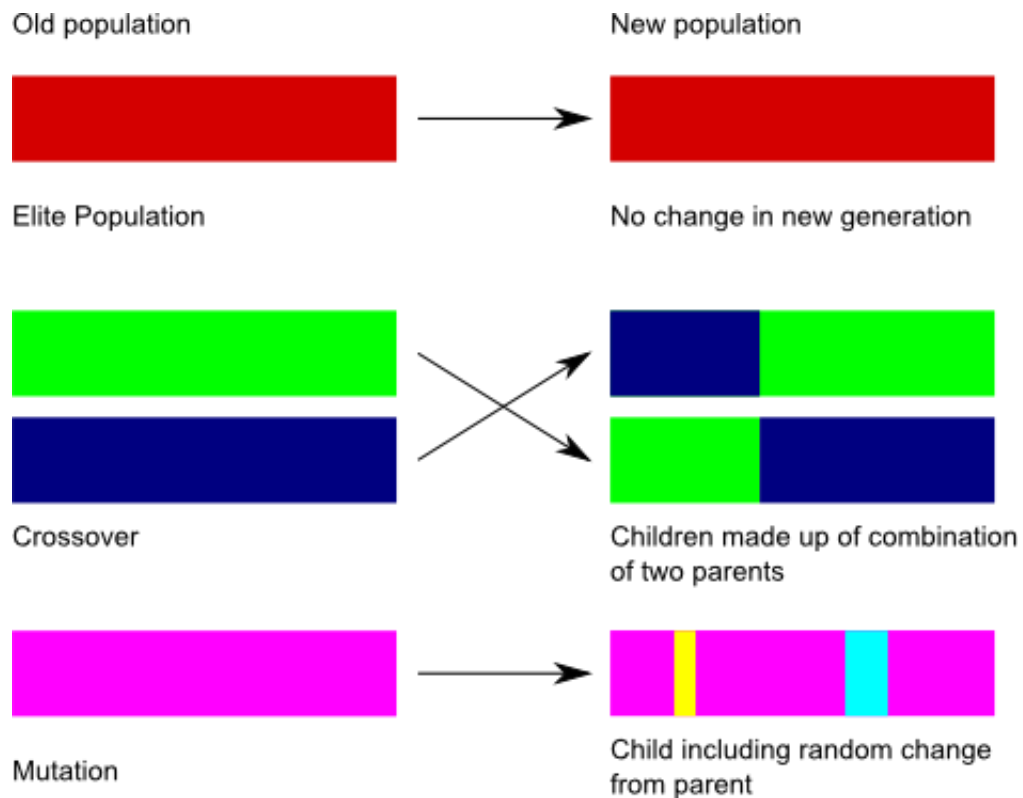


Figure 12: Diagram showing how the three different methods of reproduction occur in GA. The elite population is the best performing solutions and are unchanged between old and new populations. Crossover is method of reproduction that occurs most commonly. New solutions are made of a combination of results of old solutions. Mutation occurs to ensure there are no solutions that are unable to be reached through other processes. New solutions produced by this technique are the same as an old solution with some random change occurring.

Crossover is a way to add in variation to the new population. Strings are 'mated' to create a new population by randomly selecting which bits of two strings to switch round to each other. The probability that a string is involved in this process is determined by its fitness function.

Mutation also occurs but plays a secondary role to crossover in providing a variation in the new population. During mutation part of a string in the population is randomly changed to a new value. This is done since there are some configurations of solutions that may be otherwise lost if crossover is the only way new populations are generated. By randomly changing some solutions, GA can ensure all possible solutions can be reached.

The convergence of the population in GA is a serious problem. This is usually solved by finishing the process when improvement has fallen below a certain threshold over a defined number of generations. A maximum number of generations can also be used to stop the process.

The starting population can either be set at random or, if some information is known about where the solution is likely to be then this can be used to set a starting population which is close to the desired solution.

2.5.3 Previous Use of Genetic Algorithms

Genetic algorithms have been used for many different applications in the field of image processing. These include use in image enhancement, where they have been used to construct a fuzzy image filter [106]. A number of groups have used genetic algorithms for image contrast enhancement [107] [108] and many other image enhancement techniques. They have also been used for image segmentation [109] [110], solving many different problems within that field. Since it is wished to use genetic algorithms to solve a registration problem it is appropriate to discuss this in more detail.

2.5.3.1 Image Registration

Image registration is the alignment of images of the same object taken at either different times or using different methods. One area where GA has been used extensively is looking at images of the brain. Here computed topography (CT) scans can give good information on the hard tissue such as bone, and MR imaging provides good information on the location of soft tissues such as the brain. One group have used a method based on genetic algorithms to register these two images to each other [111]. They developed a method that allowed for more translation but also distortion of the image. A simple global elastic transform was used to model MR image distortions. This was defined by

$$\begin{aligned}
 p &= (x, y, z)^T & p' &= (x', y', z')^T = T(p) \\
 p' &= \sum_{i=0}^1 \sum_{j=0}^1 \sum_{k=0}^1 (a_{i,j,k}, b_{i,j,k}, c_{i,j,k})^T x^i y^j z^k
 \end{aligned} \tag{44}$$

where $a_{i,j,k}$, $b_{i,j,k}$ and $c_{i,j,k}$ are coefficients of the transformation. The aim of the algorithm is to find pairs of homologous points between the two images to achieve a global warping calculated with 24 coefficients.

In order to achieve a good registration a fitness function must be selected which is capable of both being fast and a good fit to the required result. A feature based fitness function was calculated for this. The air-skin interface can be extracted from both CT and MRI images. The distance between these two surfaces was used as the fitting function. Let S_i be the surface extracted from F_i and $d_E(p_1, p_2)$ be the

distance between points p_1, p_2 . The distance d_a with respect to transformation T can be defined as

$$d_a(S_1, S_2|T) = \frac{1}{\text{card } S_1} \sum_{p_1 \in S_1} \min_{p_2 \in S_2} d_E(T(p_1), p_2) \quad (45)$$

In order to save computational time a distance map D_2 can be created where each point in D_2 is equivalent to $\min_{p_2 \in S_2} d_E(T(p_1), p_2)$. In order to further increase speed a series of n points are chosen randomly from the set S_1 and the distance is computed with this new set, giving

$$d_a(S_1, S_2|T) = \frac{1}{n} \sum_{p_1 \in S_1^n} D_2(T(p_1)). \quad (46)$$

The need to do this comes from the large number of times this must be calculated when finding the minimum using genetic algorithms.

The group then used a GA to optimise this function. They found that the GA could easily get them to a location near the local minimum they wished to find, but it struggled to find the optimum solution. They therefore used a multistep process where the GA was used to find the local region the optimum solution was found in, and then a fine tuning process was carried out to improve the solution.

Since their work many other groups have looked at using genetic algorithms for image registration. These have included different ways to generate the fitness function as well as variations on the genetic algorithm itself [112]. One alternative way to generate a fitness function is to use mutual information [113].

Mutual information uses a concept from information theory to measure how similar two images are. When a maximum of mutual information is found it is hoped this will correspond to the correct registration. Mutual information of two random variables A, B with marginal probability distributions $p_A(a)$ and $p_B(b)$ joint probability distribution $p_{AB}(a, b)$ is defined as

$$I(A, B) = \sum_{a,b} p_{AB}(a, b) \log \frac{p_{AB}(a, b)}{p_A(a) \cdot p_B(b)} \quad (47)$$

This is related to the notion of entropy from information theory by the following equation

$$I(A, B) = H(A) + H(B) - H(A, B) \quad (48)$$

where $H(A)$ is the entropy of A and $H(A,B)$ is their joint entropy. Entropy is defined as

$$H(A) = - \sum_a \log p_{AB}(a, b) \quad (49)$$

It has been show that mutual information is always non-negative.

In order to find the correct registration the maximum value of $I(A,B)$ is found. This was done using GA. The main advantage of using mutual information combined with GA is that it is fully automatic and does not require any prior segmentation steps. It is also mathematically simple which helped with the early adoption of this technique.

2.5.4 Variations on GA

One group attempting to achieve surface registration used a slight variation on GA [114]. They were using an energy function based on minimising the Euclidian distance between the surface they were registering and a model surface. They allowed the surface to have translation and rotation on 3 different axes and also to scale. In order to get to a solution more quickly they used a dynamic mutation step, initially allowing the parameters to take a large range of values and so allow for mutation to cover anywhere in this region. When a solution was reached they then repeated the GA but this time looked at a smaller region. This meant that mutation occurred over a smaller region and is analogous to using a smaller step size to search carefully in a particular region. This was done repeatedly until using a smaller region no longer produced any improvement in registration.

2.5.4.1 Using Real Valued GA

When GA was first formulated, the parameters that the mutation and crossover operations were applied to were in the form of binary strings. This was done since it more closely follows ideas from genetics and allows for mutation to happen in a random way. However there are a number of disadvantages to this. Often there is some prior knowledge about the solution space so each parameter will have limits placed on it. When applying crossover or mutation to binary strings there is no requirement for the new value to still be a valid parameter, therefore a potentially computationally costly check of values must be carried out before accepting anew value as valid. An alternative way to implement this is to use the actual values of the parameters directly in the model. This removes the need to convert binary strings to parameters when calculating the cost function, resulting in a faster programme [115].

2.5.5 Using Zernike Polynomials to Model the Cornea

Accurate modelling and measurement of corneal surfaces has long been an important problem for a number of different reasons. Refractive surgery requires accurate modelling of the corneal shape to ensure a good outcome of surgery. Contact lens design and fitting is also based on corneal topography. Corneal data from a number of different imaging sources can be given as a functional representation in terms of a Zernike polynomial expression [116]. The anterior surface can be modelled by a finite series of Zernike polynomials [117]

$$C(\rho, \theta) = \sum_{p=1}^P a_p Z_p(\rho, \theta) + \epsilon \quad (50)$$

where $C(\rho, \theta)$ is the corneal surface, p a polynomial ordering index, $Z_p(\rho, \theta)$ the p th order Zernike polynomial, a_p is a coefficient associated with polynomial p , ρ, θ is the normalised position in polar coordinates and ϵ is noise due to measurement and modelling error.

The p th order Zernike polynomial is given by

$$Z_p(\rho, \theta) = \begin{cases} \sqrt{2(n+1)} R_n^m(\rho) \cos(m\theta), & \text{even } p, m \neq 0 \\ \sqrt{2(n+1)} R_n^m(\rho) \sin(m\theta), & \text{odd } p, m \neq 0 \\ \sqrt{n+1} R_n^0(\rho), & m = 0 \end{cases} \quad (51)$$

where n is the radial degree, m is the azimuthal frequency, and

$$R_n^m(\rho) = \sum_{s=0}^{(n-m)/2} \frac{(-1)^s (n-s)!}{s! \left(\frac{n+m}{2} - s\right)! \left(\frac{n-m}{2} - s\right)!} \rho^{n-2s}. \quad (52)$$

The radial degree and azimuthal frequency are integers that satisfy $m \leq n$ and $n - |m| = \text{even}$. The optimal number of Zernike terms varies depending on the topography of the eye. For normal and astigmatic corneas one study found the optimal number of terms was 11 [118]. They found that for distorted corneas the optimum number of terms varied from subject to subject. Using too many terms results in over parameterisation of the model.

2.5.6 Relevance to Our Work

Genetic algorithms have been shown to be a technique capable of finding a global minimum of an energy function. They have been extensively used in many different image processing techniques. They are therefore a likely candidate for use in our work. The problem to be discussed here is attempting to solve is the registration of a set of 3D images. This differs from traditional registration problems in that there is not any considerable area of overlap between the different images. This means

techniques such as maximum mutual information registration techniques are not any use to us. Using genetic algorithms is not limited to any particular energy function. It may also be sensible to include some fine tuning after the GA has completed, since previous groups have found it struggles to reach a final minimum but is easily able to find the region where the global minimum is located.

Zernike polynomials have been used extensively previously to describe the shape of the cornea. This indicates they could be a good choice to use to model the shape of the cornea for the purpose of registration.

2.6 Conclusion

This chapter has reviewed previous work in several different areas. Different anterior segment imaging techniques have been discussed. OCT is a promising technique for work on the human eye, the transparent nature of the eye making it well suited to the technique.

Different segmentation techniques have been examined. This has included looking at previous segmentation of anterior segment images as well as more generally discussing a number of techniques that can be used to develop segmentation techniques. A number of different useful techniques for image segmentation have been discussed. Implementing a previous group's work segmenting the anterior segment using a threshold [63] is an obvious starting point for our own work. There has also been work using level set, graph cut and graph theory techniques on ocular OCT images indicating these methods might offer promising results.

Previous 3D alignment techniques based on the retina have often used two imaging modalities to achieve segmentation. This would not work for anterior segment images as there are no features, such as blood vessels which are used in the retina, that extend across the cornea and are easily visible in many images. GA based techniques have been extensively used for different applications so should be investigated for use in alignment of 3D images.

3 Image Segmentation in 2D

3.1 Introduction

The first stage of the PhD was to develop automated 2D segmentation techniques capable of automatically detecting the cornea. In order to tackle this problem a number of different strategies were investigated. The first step was to implement an existing technique, Shen et al.'s [63] technique using a threshold to segment the images was the most promising technique found in the literature so was chosen as a starting point. Following on from this a number of different novel techniques were developed. These are split into three categories for ease of description and evaluation. The first new technique developed was a level set with shape technique based on Chan and Vese's [70] work with the addition of a shape term. The next set of techniques was based on a graph cut implementation of the level set with shape technique. Finally a graph theory based technique was developed.

The segmentation techniques were mainly implemented using Matlab. This was chosen for its ease of use. The graph theory method was implemented in C++ since this was found to be the most accurate technique and code is generally faster when implemented in C++.

An example image is shown in Figure 13. Areas of lower signal to noise ratio can be seen either side of the centre. It is difficulties encountered in identifying the posterior surface in this area that mean it is valuable to use a shape prior term for in the segmentation.




Figure 13: An example OCT image of the anterior segment.

A variety of new techniques for segmentation were developed. The first was a level set based technique. Level set techniques were chosen since it was thought they represented a good chance to develop a novel method that was successful at

segmentation. They had previously been used successfully for segmentation of retinal OCT images. Here groups had used shape terms to improve segmentation and it was hoped to replicate their success on a different part of the eye. A number of variations of the technique were tried.

Next, segmentation using graph cut technique was carried out. Graph cut techniques were investigated because when attempts were made to extend the 2D level set technique to 3D the technique proved to be too slow. Graph cuts had previously been used to increase the speed of segmentation where level sets had been previously employed [87]. It was hoped that they could be used to similarly improve speed of the level set segmentation technique developed here.

The final 2D segmentation technique to be developed was graph theory segmentation. This technique was investigated as they had been previously used successfully to segment retinal OCT images. There had also been some limited work using these to segment parts of corneal OCT images [97]. It was hoped that incorporating our shape based technique could improve the results of this technique.

These methods will be discussed in turn. A number of variations were attempted for each method. The results of these variations are compared. The best versions of each of the methods are compared to see which method performs best. Prior to the discussion of segmentation techniques there will be a brief discussion of noise removal.

3.2 Dealing with speckle noise in OCT image

OCT images all contain speckle noise which is inherent in the imaging technique. There are number of techniques that can be used to counteract the effect of this noise on the segmentation result. The most widely used method to do this is to use a median filter. This is a simple technique where each pixel is set to the median value of pixels in a group. This technique preserves edges in the image.

Another strategy for dealing with speckle noise is to alter the imaging technique. During imaging the value for each A scan can be recorded multiple times and an average value taken. Imaging speed is important when using this technique since repeat measurements must be taken fast enough to avoid a movement in the object of interest [119]. This technique has been previously used successfully but was not available to use for this study since it requires modification of the camera.

The use of a median filter as part of the pre-processing technique was tested but it did not produce significantly different results for most of the techniques. It was decided not to carry out any de-noising steps but instead to generate segmentation techniques that are robust to small amount of noise.

3.3 Segmenting Images Using a Threshold

Before developing any novel techniques the method described by Shen et al [63] was implemented. The first step of the segmentation algorithm was to apply a threshold to the grayscale image. The threshold value was calculated using Otsu's method [61], which is implemented in Matlab. This method was shown to be able to only segment part of the cornea as shown in Figure 14. The areas of lower signal to noise ratio in the middle do not appear.




Figure 14: Image segmented using basic threshold technique. Here all the pixels over a certain value have been labelled as part of the object of interest.

One way to improve the segmentation results was to use texture rather than intensity information. Matlab has implementations of different ways to do this as part of the image processing toolbox. An entropy filter was used since this has successfully been used by Shen et al., as well as others [120]. Once the entropy filter was applied a threshold was used to segment the image. Attempting this did improve image segmentation, however it was still unable to entirely pick up regions of lower signal to noise ratio. The image after the filter has been applied is show in Figure 15 and the results of the segmentation are shown in Figure 16.




Figure 15: Image after entropy filter has been applied. This is the same image that is show in Figure 13.



Figure 16: Image segmented by applying entropy filter and then a threshold.

Shen et al. [63] uses binary morphology and gradient information to improve their segmentation. These steps were also carried out during the segmentation. These steps did not lead to an improvement in the segmentation results. This is most probably due to the initial results being too far from true segmentation.

Overall this method was found to be capable of locating the anterior surface of the cornea, which is all Shen et al were attempting to do, but was unable to accurately show the location of the posterior surface. It was concluded that this method could be used to create a reasonably accurate initial guess but a different method needed to be used to improve on this.

3.4 Common Segmentation Framework

All other 2D segmentation techniques were carried out within a shared segmentation framework. This contains some pre-processing steps that are common to all different methods. Figure 17 shows the common framework for the different methods. The first step carried out is a pre-processing step. This removes

noise as well as unwanted parts of the images such as eyelids or the iris. The second step is the main segmentation step. This step changes depending on the method used. Four different categories of technique are discussed for this step. The final step is post-processing. The step consists of smoothing the acquired boundaries by either fitting a polynomial to the result or using a smoothing function.

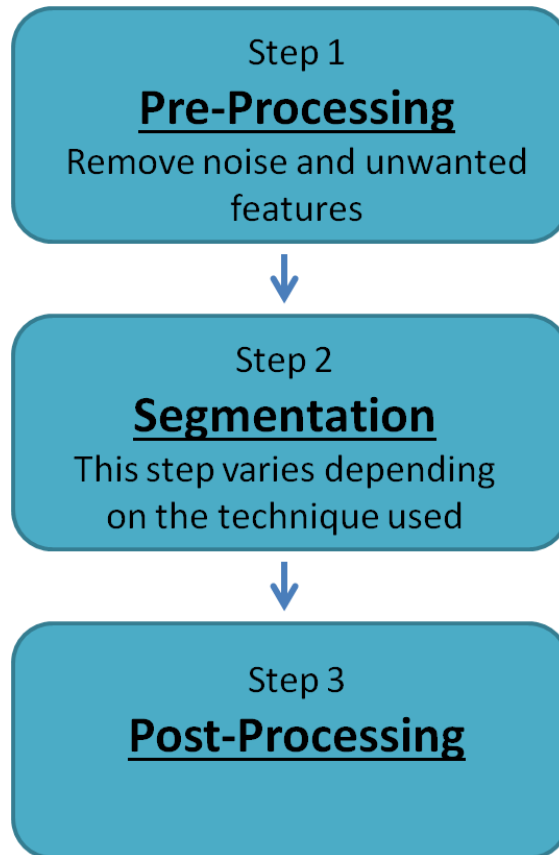


Figure 17: Generic flow chart showing the common steps of all the different segmentation techniques.

The next section will describe the pre-processing steps used. The following four sections will each be devoted to a different segmentation technique. Following this, there will be a section on post-processing.

3.5 Pre-Processing Steps

In order to improve the segmentation results a number of pre-processing steps were used. These techniques are common to the different segmentation methods so will be presented here separately. There are some slight differences in steps used for the different segmentation techniques. Which steps are used with each method is discussed when each segmentation method is introduced.

3.5.1 Removal of Iris and Central Noise Artefact

Common to all anterior segment OCT images is the central noise artefact caused by specula reflection at the apex of the cornea. Anterior segment images all contain the iris as well as the cornea. The aim of this technique is to identify the cornea so the iris can be removed to improve performance without losing important information.

The central noise artefact is characterised by having a higher signal than the rest of the image. In order to detect the central noise artefact, the mean of each column in the image was calculated. The maximum of this is always located within the central noise artefact. To remove the central noise artefact, columns with means within 90% of maximum were all set to zero. This value was used as the brightest columns were always part of the central noise, it is a conservative threshold of the noise to avoid removing parts of the image that were wanted. A small region either side of this region is also removed since these also generally contain part of the central noise artefact. There is too much noise to be able to accurately detect surfaces of interest in this region. This method has been used in previous studies [97].

The iris was located in a similar manner to the central noise artefact. It is an approximately horizontally orientated structure so all the mean value of each row in the image was calculated. The iris is characterised as a bright horizontal structure in the image behind the cornea. The rows within 90% of the maximum mean were again detected. Sometimes another area of high signal in front of the cornea was also detected. In order to only remove the iris, the program checked for a large gap in rows reaching the 90% threshold and only selected rows from below this gap if it was detected. This prevented the incorrect detection of bright rows corresponding to the apex of the cornea. The image was then cropped to remove all points below the top of the iris. The algorithms for noise and iris removal worked for all the images, although complete iris removal was not always achieved.




Figure 18: Image with iris and central noise artefact removed

3.5.2 Rotation Error

In the first data set of OCT images all the images were well aligned with the closest point to the camera being the furthest forward part of the cornea and the iris being perpendicular to the incoming light. This was not the case for the second data set. In these images there seemed to be a slightly different angle for each image. This was best seen by looking at the angle of the iris. The different angles created some problems for the segmentation method used since it makes assumptions of the corneal shape based upon the fact that the images were taken perpendicular to the plane of the cornea. Figure 19 below shows an example of an image taken at an angle.




Figure 19: Example image of cornea showing substantial deviation from perpendicular angle. Looking at the iris it can be seen the image is at an angle. Note the line down the middle is an imaging artefact so will always appear vertical in the image.

In order to solve this problem a pre-processing step that was able to correctly align the images was developed. The easiest way to identify the angle the image was taken is to look at the iris. If the alignment of the image is correct then the iris should be a horizontal structure across the image.

A method, described above, had already been developed to detect the iris based on the fact the iris is the only horizontal structure in the image, and so if the mean intensity of each row in the image is taken the row with the maximum will correspond to the iris. A simple algorithm based on this was used to rotate the image. It was assumed that rotation by an angle of less than 10° in either direction would find the correct alignment for the image. The image was then rotated at 1 degree intervals and the maximum intensity of rows calculated after each rotation.

This gave 21 possible orientations for the image. The angle with the highest maximum row intensity was assumed to be the correct orientation since that corresponded to having a horizontal iris. The algorithm worked successfully with all images in the data set.

Medical imaging on large scale applications can often suffer from images being taken that are lower quality than could be achieved with the imaging equipment under perfect conditions.

3.5.3 Eyelid Removal

Some of the images of the cornea also include parts of the eyelids in the image. This is a problem since it interferes with the identification of the cornea. It is therefore desirable to develop a method that can be used to remove the eyelid from images prior to segmentation. The angle at which the OCT images are taken at determines the chance of a section of the eyelid being included in the image. Non horizontal angles, such as those used when acquiring 3D information by taking a series of radial scans, are more likely to have parts of the eyelid included.

Three different techniques were attempted; all based the same initial premise. A threshold technique was used to segment the image. The anterior surface can be detected reasonably accurately by this technique. This surface would include both cornea and the eyelid. The eyelid can then be detected and removed by studying this line. Figure 20 shows an example image containing parts of the eyelid.




Figure 20: Example image showing image including eyelids. Note where eyelid is present it blocks out all lower features in the eye so parts of the iris are missing in this picture.

3.5.3.1 *Initial Segmentation*

The first step of eyelid removal is to segment the images. This is done using Otsu's method [61] to select a suitable threshold to segment the image. Once the image had been segmented, a diamond shaped structuring element was used to dilate the segmented image. This was done as the threshold segmentation method consistently underestimates the area of interest in the image. Dilating the image also causes different sections the cornea to be joined together. The largest segmented object is then found and assumed to be the cornea and parts joining to it. A check is carried out on the second largest object. If it is similar in size of the largest object it is also assumed to also be an object of interest and included. This was done to ensure that objects that are noise artefacts are rejected but those that are part of the cornea or eyelid are included. The regions are joined together using straight lines and the boundary of this object is found. It is assumed that this object contains all of the cornea and possibly including eyelids if they are in the image. Figure 21 shows the example image with this boundary marked on it.

The anterior surface of this shape can be found by taking the first set of boundary points from the start until the furthest right edge is reached. Using this surface, three different methods were attempted to discover the location of the boundary with the eyelid. These were a gradient based method, a method based on finding on peaks, and finally a method based on finding the maximum z value of the boundary, where z is the vertical axis of the image and is valued at zero at the top of the image.




Figure 21: Example image with the anterior surface detected using the method described above. The top surface has been detected includes both the cornea which we want and eyelids which we wish to remove. The challenge is accurately finding this boundary.

3.5.3.2 *Max z value of boundary*

The most successful method was based upon finding the max z value of the boundary. The image was split into two halves at the apex of the cornea. This was determined by finding the minimum in the central region and assuming this corresponds to the location of the centre of the cornea.

The maximum value of the anterior surface was then found in each half of the image. The location of this corresponds to the join of the eyelid and the cornea. The image was then cropped at this location. If there was no eyelid present in the image, the maximum point instead corresponds to either the end of the cornea or the edge of the image. In either of these cases, cropping the image here does not result in the loss of any useful parts of the cornea. This method proved the most successful of the three different methods that were tried, so was used in all future segmentation. The result of this method on the example image is show in **Error! Reference source not found.**




Figure 22: Example image after eyelid has been removed

3.5.3.3 *Other methods*

Rather than looking for a global maximum to find the boundary of eyelid methods looking for local minimum were also attempted. Two variations of this method were used one looking at the gradient of the boundary the other looking for local peaks. The same problem was encountered with both techniques. The middle region of the cornea has low signal to noise ratio. It was not always possible to accurately detect the anterior surface in this region. This caused gaps and extra turning points in the gradient of the anterior surface in some images,. causing this method to fail at accurately detecting the eyelid

3.6 Level Set Method

The first technique developed used level set based methods to segment the images. Level set methods were chosen because several groups have previously used this technique to successfully segment medical images. The most relevant to this work is a study by a group that used a level set formulation with a shape constraint to segment OCT images of rodent retina [79].

The first attempt made was to use one of the simplest and most popular level set methods available: Chan and Vese's [70] active contour without shape. This method attempts to segment an image into two regions based on intensity. The aim is to minimise the following energy expression

$$E(\Phi) = \lambda_1 \int_{\Omega} (I(x,y) - u)^2 H_{\epsilon}(\Phi) + (I(x,y) - v)^2 (1 - H_{\epsilon}(\Phi)) dx dy + \lambda_2 \int_{\Omega} \delta_{\epsilon}(\Phi) \nabla \Phi dx dy \quad (53)$$

where $I(x,y)$ is image intensity, u is mean intensity inside the curve, v is mean intensity outside of the curve, $H_{\epsilon}(\Phi)$ is the Heaviside function and λ_1 and λ_2 are constants that determine the weighting of each term. u and v can be defined as

$$u(\Phi) = \frac{\int_{\Omega} I(x,y) H_{\epsilon}(\Phi) dx dy}{\int_{\Omega} H_{\epsilon}(\Phi) dx dy} \quad (54)$$

$$v(\Phi) = \frac{\int_{\Omega} I(x,y) (1 - H_{\epsilon}(\Phi)) dx dy}{\int_{\Omega} 1 - H_{\epsilon}(\Phi) dx dy} \quad (55)$$

This method has been previously implemented in Matlab as part of the *creasag* program [121]. An example image segmented using this method is shown in Figure 23. This technique still struggles to locate the low signal to noise ratio regions the threshold method also struggled to find.



Figure 23: Image segmented using “Active contour without edges” method

This result confirmed it is necessary to incorporate shape information into the model to obtain an accurate segmentation.

The initial work, shown in the example image above, was carried out without any pre-processing steps. When it is compared with other techniques, all the pre-processing techniques described in section 3.5 were used.

3.7 Level Set with Shape

In order to improve the previous method, a shape term can be used to help locate parts of the image with a low signal to noise ratio. For this technique all the pre-processing techniques described in section 3.5 were used. Following the pre-processing, the first step of the segmentation was to make an initial estimate of segmentation using the thresholding technique described above. This was used to create a shape term. Finally, a level set function was evolved to improve the contour. The steps will be discussed in turn.

3.7.1 Initial Estimate of Segmentation

The initial investigation into the use of techniques based upon applying a threshold to an image shows that it is capable of finding a good estimate of the anterior surface of the cornea. Following on from the initial pre-processing step an entropy filter was applied to the image to obtain texture information. This texture image was then segmented using a threshold, and the results assumed to be a good initial estimate of the location of anterior surface of the cornea.

This initial estimate was used to initialise the level set function and provide an initial estimate for the shape term.

3.7.2 Evolution of Curve Using a Level Set Function with a Shape Constraint

In the level set technique, the active contour that is used to segment the image is represented by the zero level set of a function. This function can then be evolved to move the contour in an image. In order to control the evolution of the curve an energy function for the image must be created. The method used here contains three terms

$$J(\Phi) = J_1(\Phi) + J_2(\Phi) + J_3(\Phi) \quad (56)$$

J_1 is a region based term, J_2 a contour smoothing term and J_3 a shape based term, each term is discussed in detail in the next section. This method is similar to that used by Yazdanpanah et al when attempting to segment the retina into multiple regions as well as work looking at the segmentation of cardiac MR images [79, 81].

3.7.2.1 Model Matching – a Regional Based Term

The regional based comes from the Chan-Vese model of active contour without edges and attempts to split the image into two regions based on them having similar regional statistics [70]. The energy term corresponding to this is

$$J_1(\Phi) = \lambda_1 \int_{\Omega} (I(x,y) - u)^2 H_{\epsilon}(\Phi) + (I(x,y) - v)^2 (1 - H_{\epsilon}(\Phi)) dx dy \quad (57)$$

where $I(x,y)$ is image intensity, u is mean intensity inside the curve, v is mean intensity outside of the curve, $H_{\epsilon}(\Phi)$ is the Heaviside function and λ_1 is a constant that determines weighting of this term.

3.7.2.2 Contour Smoothing Term

The contour smoothing term is a function of the length of the contour and acts to smooth the curve.

$$J_2(\Phi) = \lambda_2 \int_{\Omega} \delta_{\epsilon}(\Phi) \nabla \Phi dx dy \quad (58)$$

where $\delta_{\epsilon}(\Phi)$ is an approximated delta function corresponding to gradient of Heaviside function, and λ_2 is a constant determining weighting of this term.

3.7.2.3 Shape Prior Term

To improve the segmentation to cover areas with lower difference in intensity a shape prior term is introduced. The surface of the cornea is approximately elliptical so an elliptical shape was used.

An ellipse was fitted to the anterior surface of the curve using a least squares fitting method [122]. The distance between a point and the ellipse is then defined as

$$D = ax^2 + by^2 + cxy + dx + ey + f \quad (59)$$

where a, b, c, d, e, and f are parameters that define an ellipse. A second ellipse was then generated by adjusting the ellipse parameters to give a second ellipse with the same centre but smaller major and minor radius. The two different D matrices obtained were then multiplied together to create a single matrix

$$D = D_{upper} \cdot D_{lower} \quad (60)$$

The D matrix has values everywhere on the image and is zero on two curves corresponding to the two ellipses the surface is fit to.

The shape constraint used was given by

$$J_3(\Phi) = \lambda_3 \int_{\Omega} D^2 \delta_{\epsilon}(\Phi) \nabla \Phi dx dy \quad (61)$$

Where D is the shape prior term, $\delta_{\epsilon}(\Phi)$ is a delta function corresponding to gradient of Heaviside function, and λ_3 is a constant determining weighting of this term.

3.7.2.4 **Minimisation of the Energy Function**

The above terms are all substituted into the original expression and our energy to be minimised is defined as

$$J(\Phi) = \int_{\Omega} \lambda_1 [(I(x, y) - u)^2 H_{\epsilon}(\Phi) + (I(x, y) - v)^2 (1 - H_{\epsilon}(\Phi))] + A(x, y) \delta_{\epsilon}(\Phi) \nabla \Phi dx dy \quad (62)$$

where $A(x, y) = \lambda_2 + \lambda_3 D^2$.

To evolve the contour and find the best fit for the curve on the image the energy function calculated above must be minimised. This is done in a series of steps.

Firstly, holding Φ constant, the intensity inside and outside the curve (u and v) is calculated.

The next step is to calculate the value of D while holding Φ constant. It is usual to calculate the shape constraint each iteration of the code updating the expression each time to get a better fit of results. This approach was not used here. Instead a

good initial estimate of the upper surface of the cornea was obtained. The shape term was then updated every 30 steps, generating a new shape function at that point. The main reason for doing this was because it resulted in a substantial improvement in the running speed of the program.

Once the intensities and distance function are calculated, these are both held fixed and the Euler-Lagrange equation is used calculate the parameterised descent. From the Calculus of Variation [123] it can be shown that the 2D scalar function $\Phi(x, y)$ that minimises the energy function in equation 62 is the one that solves the PDE

$$\frac{\partial J}{\partial \Phi} - \frac{d}{dx} \left(\frac{\partial J}{\partial \Phi_x} \right) - \frac{d}{dy} \left(\frac{\partial J}{\partial \Phi_y} \right) = 0 \quad (63)$$

The partial derivative of J in with respect to Φ , Φ_x and Φ_y .

$$\frac{\partial J}{\partial \Phi} = \lambda_1 \frac{\partial H}{\partial \Phi} [(I - u)^2 - (I - v)^2] + A\delta'(\Phi) |\nabla \Phi| \quad (64)$$

$$\frac{\partial J}{\partial \Phi_x} = A\delta(\Phi) \frac{\Phi_x}{|\nabla \Phi|} \quad (65)$$

$$\frac{\partial J}{\partial \Phi_y} = A\delta(\Phi) \frac{\Phi_y}{|\nabla \Phi|} \quad (66)$$

where $\delta' = \partial \delta / \partial \Phi$.

Taking the derivative of equation 65 above with respect to x gives,

$$\frac{d}{dx} \left(\frac{\partial J}{\partial \Phi_x} \right) = A_x \delta(\Phi) \frac{\Phi_x}{|\nabla \Phi|} + A\delta'(\Phi) \frac{\Phi_x^2}{|\nabla \Phi|} + A\delta(\Phi) \frac{d}{dx} \left(\frac{\Phi_x}{|\nabla \Phi|} \right) \quad (67)$$

Similarly the derivative of equation 66 with respect to y is

$$\frac{d}{dy} \left(\frac{\partial J}{\partial \Phi_y} \right) = A_y \delta(\Phi) \frac{\Phi_y}{|\nabla \Phi|} + A\delta'(\Phi) \frac{\Phi_y^2}{|\nabla \Phi|} + A\delta(\Phi) \frac{d}{dy} \left(\frac{\Phi_y}{|\nabla \Phi|} \right) \quad (68)$$

Equation 62 can now be rearranged to be

$$\lambda_1 \frac{\partial H}{\partial \Phi} [(I - u)^2 - (I - v)^2] - \left[\nabla A \cdot \frac{\nabla \Phi}{|\nabla \Phi|} + A \text{div} \left(\frac{\nabla \Phi}{|\nabla \Phi|} \right) \right] \delta(\Phi) = 0 \quad (69)$$

This can then be used to evolve Φ using an artificial time t . The following update equation is then found

$$\frac{\partial \Phi}{\partial t} = -\lambda_1 \frac{\partial H}{\partial \Phi} [(I - u)^2 - (I - v)^2] + \left[\nabla A \cdot \frac{\nabla \Phi}{|\nabla \Phi|} + \text{Adiv} \left(\frac{\nabla \Phi}{|\nabla \Phi|} \right) \right] \delta(\Phi) \quad (70)$$

Using this method did not result in good segmentation of the images. An example result is shown in Figure 24. The poor results were due to the shape term not working as hoped to fill in gaps of the segmentation that were found when using only intensity information. In order to improve the results a different approach to formulating the shape constraint was used. This is described in the next section.



Figure 24: Image segmented using the first attempt at level set function with shape

3.7.2.5 Alternative Methods for Shape Term

There are many different ways to implement a shape term in a level set function. These generally involve creating a signed distance function that corresponds to the shape that is being segmented [124].

One method used by Cremers et al. [125] was to use both prior shapes and dynamic labelling. Letting Φ be the level set function for segmentation and Φ_0 a signed distance function for the given shape, a shape difference function was created as follows

$$J_{shape}(\Phi) = \int_{\Omega} (\Phi - \Phi_0)^2 dx dy \quad (71)$$

The problem with this integral is that the shape prior is enforced over the whole domain. This could create problems due to background objects affecting the segmentation of an image. Another function L can be introduced into the integral to deal with this problem [126].

$$J_{shape}(\Phi) = \int_{\Omega} (\Phi - \Phi_0)^2 (L + 1)^2 dx dy \quad (72)$$

where L defines where the shape prior should be active in the image domain. Regions where $L = -1$ are excluded from the integral. How the labelling region is defined is an important problem when using this method. The simplest way to do

this is to predefine the labelling region in advance. This has the disadvantage of requiring some knowledge of where the object to be segmented is located. In this application this is not a major problem since thresholding can be used to get an approximate location of the cornea. This can be used to restrict the shape term to regions near the approximate cornea location.

3.7.2.6 **New Evolution Equation**

If we substitute equation 72 into equation 62 then a new evolution equation can be derived. The energy function to be minimised now becomes

$$\begin{aligned}
& J(\Phi) \\
& = \int_{\Omega} \lambda_1 [(I(x, y) - u)^2 H_{\epsilon}(\Phi) + (I(x, y) - v)^2 (1 - H_{\epsilon}(\Phi))] + \lambda_2 \delta_{\epsilon}(\Phi) |\nabla \Phi| \\
& \quad + \lambda_3 (\Phi(x, y) - \Phi_0(x, y))^2 (L + 1)^2 dx dy
\end{aligned} \tag{73}$$

Based on the Euler – Lagrange equation as before

$$\frac{\partial J}{\partial \Phi} - \frac{d}{dx} \left(\frac{\partial J}{\partial \Phi_x} \right) - \frac{d}{dy} \left(\frac{\partial J}{\partial \Phi_y} \right) = 0 \tag{74}$$

Equation 70 now becomes

$$\frac{\partial J}{\partial \Phi} = \lambda_1 \frac{\partial H}{\partial \Phi} [(I - u)^2 - (I - v)^2] + \lambda_2 \delta'(\Phi) |\nabla \Phi| + 2\lambda_3 (L + 1)^2 (\Phi - \Phi_0) \tag{75}$$

The other derivatives are unaffected, other than setting $A = \lambda_2$ making it no longer a function of (x,y). This gives the gradient descent equation as

$$\begin{aligned}
\frac{\partial \Phi}{\partial t} = & -\lambda_1 \frac{\partial H}{\partial \Phi} [(I - u)^2 - (I - v)^2] - 2\lambda_3 (L + 1)^2 (\Phi - \Phi_0) \\
& + \lambda_2 \operatorname{div} \left(\frac{\nabla \Phi}{|\nabla \Phi|} \right) \delta(\Phi)
\end{aligned} \tag{76}$$

3.7.2.7 **Defining the Shape Constraint**

One of the problems faced when using the above method is how the shape term is constructed. The aim is to generate a shape that is as close as possible to the shape of the cornea and then use this to influence the contour to follow the edges of this shape.

The initial technique attempted was to fit an ellipse to the anterior surface and then generate a second ellipse with a smaller major and minor axis and use this to approximate the bottom surface. This method was initially successful, however

when updating the shape expression as the curve developed problems arose. The method being used did not lead to a fixed sign for the new shape. This prevented a stationary end point from being reached.

To solve this problem a new method of generating a smaller ellipse was used. The ellipse fitted to the anterior surface was used to generate a signed distance function (SDF). The central width of the cornea was manually measured on a few example images. The average of these values was used to estimate the bottom edge of the cornea at the centre. The value of the SDF at this point was then subtracted from the original SDF to set a new zero point here. This method successfully generated a new ellipse that was parallel to the original one. However this is not a good approximation to the cornea. The cornea thickness increases from the centre towards the edges of the image.

In order to attempt to mimic this effect initially the minor axis had been decreased twice as much as the major axis. In this case the amount the SDF was varied was based on horizontal position in the image. A quadratic relationship was used with the amount each row was adjusted by being given by the following expression

$$SDF2(x, y) = SDF1(x, y) - SDF1(x_t, y_t) - c * (x - x_t)^2 \quad (77)$$

Where SDF1 and SDF2 are the initial and altered signed distance functions respectively, (x_t, y_t) is a point 30 pixels below the anterior ellipse, and c is a constant obtained empirically. An example of the shape image is shown in Figure 25 below. Colour has been used to make it clearer.



Figure 25: Example of the shape constraint used to guide segmentation. Here dark blue corresponds to zero and below and red to higher values.

3.7.3 Initial Results

This method produced in a successful segmentation of most of the images. The results for the same image as shown previously are shown in Figure 26 and Figure 27 below. Full results will be presented in Chapter 5.



Figure 26: Image segmented using level set with shape




Figure 27: Image segmentation using level set with shape marked on original image

One of the problems encountered was how to select the weighting of different parts of the energy expression. The strength of the shape constraint was varied for all the different images. There was no one value that worked for all images. Some images were better segmented with a very weak shape constraint, whereas other performed better with a stronger shape term. The values of λ_3 tested varied from 0.1 to 0.6. It is possible the images not successfully segmented may have worked if a wider range of values was used. In order for a fully automatic segmentation program to be used, the values of the weighting parameters must be calculated the same way for all images. Using dynamic methods where λ_3 is a function of number of iterations could improve results. Alternatively if a link between image quality and λ_3 value could be found this could be used to determine best value for each image.

3.7.4 Improvements to the Level Set Method using Dynamic Weighting

The results from the level set with shape method were promising. Therefore ways to improve the model with relatively minor changes were tested. Previous image segmentation methods have improved segmentation by varying the weighting of different terms in the energy expression. This can be done in a number of different ways.

One way that was investigated was adjusting the strength of the shape term based on image quality. The images were graded on quality by the expert manual observer when they carried out the segmentation. For each image, the value of the shape coefficient that gave the highest DSC value was found. Comparing the two results, no correlation was found between rated image quality and value of shape coefficient that gave the best result. Given the lack of any automated way of rating image quality, and the initial negative result, this line of investigation was not continued.

3.7.4.1 *Dynamic Changing of Weighting Factors*

The initial model of segmentation used constant values of the various coefficients for segmentation. Other studies have found improved results by varying the values of the coefficients by iteration number. This could have the advantage of using mainly image information to enable to correct region of the image to be detected. These methods were therefore attempted for this study. A number of variations were attempted initially based on previous work. A number of variations are described here these all produced worse quality results compared to using constant weighting.

In STACS [81] there are four terms in the energy function. The first two are regional and edge information from the image, the third term is their shape term and the final turn a smoothing term. When minimising their energy function the following expressions were used for the different coefficient

$$\lambda_1(n) = \lambda_1(1) - \frac{n[\lambda_1(1) - \lambda_1(N)]}{N} \quad (78)$$

$$\lambda_2(n) = \frac{1}{2}[\lambda_2(1) - \lambda_2(N) \left[1 + \cos\left(\frac{n\pi}{N}\right)\right] + \lambda_2(N)] \quad (79)$$

$$\lambda_3(n) = \frac{\lambda_3(N) - \lambda_3(1)}{\cosh\left[10\left(\frac{n\pi}{N} - 1\right)\right]} + \lambda_3(1) \quad (80)$$

where n is iteration number, N the total number of iterations and $\lambda_i(1)$ and $\lambda_i(N)$ the initial and final value of the coefficients. λ_4 was held constant throughout the minimisation process. These expressions correspond to slowly decreasing λ_1 and λ_2 and a λ_3 value that increases dramatically at the end of the segmentation process. The evolution of these values is shown in Figure 28 below.




Figure 28: Graph showing evolution of different coefficient in Pleumpitiyawej et al.'s study [81]

Since this method had shown an improvement in previous results, a modified version was tried. A number of different methods of changing the strength of different terms as segmentation progressed were also tried. Different strategies attempted were

$$\begin{aligned}
 \lambda_3(n) &= 0 && \text{when } n < 100 \\
 \lambda_3(n) &= \frac{\lambda_3(N)}{2} && \text{when } 100 \leq n \leq 1500 \\
 \lambda_3(n) &= \lambda_3(N) && \text{when } n > 1500
 \end{aligned} \tag{81}$$

where n is number of iterations, and N is max number of iterations (3000 for this implementation). λ_1 and λ_2 were held constant throughout the process. This method represents a step increase in strength of shape term. Results using this method were worse than when a constant λ_3 value was used.

The method used by Pluemptoyawei et al. [81] was then used. It was altered to remove the extra term giving a system of

$$\lambda_1(n) = \lambda_1(1) - \frac{n[\lambda_1(1) - \lambda_1(N)]}{N} \quad (82)$$

$$\lambda_2(n) = \lambda_2(N) \quad (83)$$

$$\lambda_3(n) = \frac{\lambda_3(N) - \lambda_3(1)}{\cosh[10(\frac{n\pi}{N} - 1)]} + \lambda_3(1) \quad (84)$$

This method again resulted in a decrease in the quality of the resulting segmentation.

The methods described so far are based upon increasing the shape coefficient with iteration. They all decrease the effectiveness of the segmentation therefore it is useful to attempt to implement a method where the strength of the shape coefficient is reduced by iteration. The following expression was used

$$\lambda_3(n) = \left(1 - \frac{n}{2N}\right) \lambda_3(N) \quad (85)$$

with λ_1 and λ_2 being held constant throughout the process. This method results in a steady increase in the value of the shape coefficient. Implementing this method also results in a decrease in effectiveness of the segmentation model.

Three different methods of altering the strength of the coefficients were attempted. These have previously been reported to give improved segmentation results. All three techniques tested for altering the strength of the shape term using iteration number produced worse results compared to using a constant strength shape constraint. There are many more ways that the segmentation could be changed with iteration number. However due to repeated negative results this work was not continued. Instead a constant strength shape coefficient was used throughout the energy minimisation process.

3.7.4.2 **Chan Vese with Shape and Edge (CVWSe)**

Another way of altering the strength of the shape coefficient is to use information from the image to vary it across each image. This has been done in previous studies [81] where edge information was used to weight the shape coefficient. In order to do this the following expression was used as an edge detector

$$g = \frac{1}{1 + k\nabla I} \quad (86)$$

where k is a constant that determines the strength of the gradient term and ∇I is the gradient of the image. g is 1 at regions where the gradient is zero and decreases at regions where there is a high gradient. A strong gradient corresponds to a clear edge, an example of what this function looks like is shown in Figure 29. In order to improve the segmentation λ_3 was replaced by $\lambda_3 * g$ in the evolution equation:

$$\frac{\partial \Phi}{\partial t} = -\lambda_1[(I - u)^2 - (I - v)^2] + \lambda_2 \operatorname{div} \left(\frac{\nabla \Phi}{|\nabla \Phi|} \right) - 2\lambda_3 * g(\Phi - \Phi_0) \quad (87)$$

This model includes the standard Chan Vese model plus the addition of shape and edge information. It will therefore be referred to as the CVWSe model when being compared to other segmentation techniques.




Figure 29: Example of gradient function g . Note the anterior surface has obvious black line but the posterior surface contains a section with no clear edge present.

3.7.4.3 *Linking Weighting of Different Coefficients with Image Intensity*

The images used in this study contained some regions with strong signal and other regions with weaker signal. It is required to detect the cornea in both regions. The shape term is used to improve the segmentation in regions with lower signal to noise ratio. In regions where the signal is stronger, the shape term is not needed to improve segmentation therefore the strength of the shape term was linked to image intensity. If the shape coefficient had an inverse dependence on image intensity then it may improve the results.

An extra term was added to the shape function

$$\lambda'_3 = \lambda_3 \frac{0.1}{I(x,y) + 0.1} \quad (88)$$

where λ'_3 is the new shape coefficient, λ_3 is a constant and $I(x,y)$ is the image intensity at x,y . 0.1 was added to the denominator to prevent large values being caused by dividing by very low values. This method will reduce the value of the shape coefficient in regions of high image intensity. An example of what this function looks like is shown in Figure 30 below. A Gaussian filter was applied to the image first to reduce the effect of any spikes in image intensity due to noise. During comparison with other methods this method will be known as CVWSiw.




Figure 30: Example of intensity weighting function used. Darker areas correspond to lighter patches on original image and the shape term will have a smaller effect in these locations. Again this function causes the difficult to detect region of the cornea to have an increase dependency on the shape term.

3.7.5 Use of Texture Information to Improve the Model

When carrying out segmentation using Chan and Vese active contours without edges there are a number of different types of image information that can be used. The methods that have been described so far are based on using image intensity to distinguish the two regions. An alternative to using intensity is to use the texture of an image. Both these methods have been extensively used in image segmentation. There are a number of different ways to represent image the texture of an image. The method used is a model matching method which assumes two different statistical models can be created one for object of interest and one for the background.

The image domain is separated into two regions by a contour C . This contour is improved by maximising the probability that u_1 and u_2 are random samples that correspond to object model M_1 and background model M_2 . The contour is updating by maximise the function

$$E_1(\Phi) = p(u|C, M_1, M_2) \quad (89)$$

where $p(u|C, M_1, M_2)$ is the joint probability density function of image intensities u given the contour C and models M_1 and M_2 .

If M_1 and M_2 are assumed to be statistically independent, it is possible to rewrite the above as a product of two probability density functions

$$E_1(\Phi) = p_1(u_1|C)p_2(u_2|C) \quad (90)$$

Taking the negative log this becomes

$$E_1(\Phi) = -\ln(p_1(u_1|C)) - \ln(p_2(u_2|C)) \quad (91)$$

Assuming that these are independent and C is a zero level set of function Φ equation 57 becomes

$$E_1(\Phi) = \int_{\Omega} -\ln[p_1(u(x, y))] H_e(\Phi(x, y)) - \ln[p_2(u(x, y))][1 - H_e(\Phi(x, y))] dx dy \quad (92)$$

where $H_e(\Phi(x, y))$ is the regularised Heaviside function representing the pixels inside the contour and $1 - H_e(\Phi(x, y))$ represents the pixels outside. Both the object and background pdfs are Gaussian with means m_1 and m_2 , and variances s_1 and s_2 . These will be calculated again each time the contour is updated. This new energy function is used in the following equation (same as equation 56 presented earlier)

$$E(\Phi) = \lambda_1 E_1(\Phi) + \lambda_2 E_2(\Phi) + \lambda_3 E_3(\Phi) \quad (93)$$

where $E_2(\Phi)$ is the curvature smoothness term and $E_3(\Phi)$ the shape term. Terms $E_2(\Phi)$ and $E_3(\Phi)$ are unchanged in this model and have been discussed in previous sections.

The negative log of the pdfs can be rewritten as

$$M_i = \frac{1}{2} \ln(2\pi s_i) + \frac{(I(x, y) - m_i)^2}{2s_i} \quad (94)$$

where s_i is the variance of pixels inside or outside contour, m_i is mean of pixels inside or outside contour and $i = 1, 2$ to specify area inside or outside contour.

This gives an evolution equation for this model of

$$\frac{\partial \Phi}{\partial t} = -\lambda_1 [M_1 - M_2] + \lambda_2 \operatorname{div} \left(\frac{\nabla \Phi}{|\nabla \Phi|} \right) - 2\lambda_3 (\Phi - \Phi_0) \quad (95)$$

The rest of the algorithm described earlier is unchanged. This model is referred to as CVWSt when being compared to other methods.

3.7.6 Extending the Segmentation to Include the Iris

So far all the methods used started with a pre-processing step which included the removal of the iris from the image. This is done because its presence in the image negatively affects the results since it deviated from the model for the shape of the cornea. The main focus for the study is the cornea however it is also desirable to segment the iris for various reasons. For example detecting the iris allows measurement of anterior chamber angle, which is important in Glaucoma investigation [127]. Following the success of the segmentation program when investigating the cornea it was decided to extend the segmentation to include the iris.

Using the detection of the location of the iris in the initial steps a sub image containing only the iris and a limited part of the sclera is created. In the images used the iris is much easier to detect than parts of the cornea. The images of the iris could be segmented and the position of the anterior surface of the iris detected using the threshold technique described previously [43] and exploited here as an initial estimate.

Specifically an entropy filter is applied to the image. This is a measure of the texture of an image and gives a better contrast between object and background than simple intensity does. The new image is then segmented using a threshold. The value used for the threshold is determined using Otsu's method [61]. A three way Otsu's method is used since there are now three regions in the image. These are; an empty padding area that was produced by the earlier rotation step to align the iris horizontally, the background of the image, and the object of interest in the image. If a two region segmentation method is used then only the difference

between the image and the empty padding area will be detected since the latter is completely uniform.

Following the initial segmentation, morphological operations were carried out to improve the segmentation. The image was eroded using a diamond shaped structural element. This was done since it was found that the segmentation was a slight over estimate of the region of interest.

Once segmentation of the region containing the iris is complete the best method of recombining the two images must be found. Simply segmenting the two images and reattaching them suffers from the problem that unless the two segmentations are identical on the boundary, a discontinuity will exist. In order to solve this problem an overlap region was used. When cutting the image to remove the iris an extra section of the image was also added so the two images overlap. When recombining the two images the boundaries in this region are combined. At the top of the region the boundary used corresponds to the top image and at the bottom the boundary used corresponds to the bottom image. In between these points the following expression is used to pick the boundary

$$bound_{new} = bound_{upper} - \frac{i}{n}(bound_{upper} - bound_{lower}) \quad (96)$$

where $bound_{new}$ is x coordinate of combined boundary, $bound_{upper}$ the x coordinate of boundary found through corneal segmentation, $bound_{lower}$ the x coordinate of boundary found through iris segmentation, n is the number of pixels in overlap region and i is the y coordinate with in the overlapping region ie 1 at top of region and n at bottom. This gives a boundary that is smooth in transition between the two regions.

Once the images have been combined an image is generated that has been segmented, including the iris. The method described here works well at detecting the anterior surface of the iris. The posterior boundary of the iris is harder to detect. This structure is not well defined on the image so this is probably a limitation in the imaging technique rather than our segmentation method. An example image of the results of this technique is shown below in Figure 31. The jagged edge on bottom right of image is mostly likely incorrect segmentation. The anterior boundary of the cornea is smooth right through to the Sclera and there is no obvious boundary

between the two different segmentation techniques. The posterior boundary of the cornea is also well defined, as is the top surface of the iris.




Figure 31: Example of segmentation result including the iris. Top section of the image is segmented using CVWSiw technique. Lower section including the iris is segmented using threshold technique described immediately above this. Well defined anterior boundary of iris can be seen. The lower surface of iris is a very rough surface indicating probable inaccurate segmentation.

3.7.7 Alternative Ways to Generate Shape Energy

The methods presented so far have used a parametric shape term, assuming that the shape of the cornea can be considered as elliptical. The results of our segmentation show that this is not an unreasonable assumption to have made. There are other ways of approaching the use of shape for image segmentation. Statistical shape priors can be used to generate a shape energy term and has widely been used for different applications in image processing.

The use of a statistical shape model has the advantage of being able to generate shapes that are not easily described by comparison to geometric shapes. The most common variations from these shapes are also identified and quantified using this model. In order to carry out segmentation using this method set of training images is needed. Once the manual annotations were acquired these could be used to train the model and attempt segmentation using this technique.

Bresson et al. [77] have made their code publically available for others to use. This was used to see if it would result in an increased performance of the segmentation model. Since only had one data set of 39 images was available each image was

segmented using the other 38 to make up the reference data set. This maximises the amount of data in the training set and allows for a full comparison to be made between the different methods. The use of this segmentation technique produced poor results so work was focused on using a parametric shape as described above.

3.8 Graph Cut Segmentation

One of the weaknesses of using level set based techniques is the speed of the segmentation. Level set is a slow segmentation technique. Work has been done previously implementing level set models using graph cuts in order to increase the speed of the segmentation. It was therefore decided to implement the CVWSiw model using graph cuts to attempt to speed up the code. This technique uses all the pre-processing techniques described in section 3.5.

In general when carrying out graph cut segmentation the energy function to be minimised should be in the following form

$$E(x) = \sum_i E_i(x_i) + \sum_{i,j \in N} E_{ij}(x_i, x_j) \quad (97)$$

where n is a neighbourhood system on pixels, $E_i(x_i)$ is a function from image that measures the cost of assigning label to pixel x_i to pixel i , and $E_{ij}(x_i, x_j)$ is a function that measures the cost of assigning labels (x_i, x_j) to adjacent pixels i, j [86]. This is discussed in more detail in literature review section the equation and is repeated here for clarity.

In order to implement the function it was necessary to create a shape term that could be used in a graph cut function. The shape mask was developed from the initial estimate of segmentation using the same method used during the level set implementation of the code, i.e. fitting an ellipse to the anterior surface then creating a second related ellipse to model lower surface. A binary shape mask was used rather than a distance function since this is similar to how shape has previously been added to graph cut segmentation.

This means that $E_i(x_i)$ in equation 99 now becomes

$$E_i(x_i) = |I(p) - c_1| + |M(p) - 1| \quad \textit{inside the contour} \quad (98)$$

$$E_i(x_i) = |I(p) - c_2| + |M(p) - 0| \quad \textit{outside the contour}$$

where c_1 and c_2 are mean intensities for inside and outside contour, $I(p)$ is image intensity at p and $M(p)$ value of mask at p . $E_{ij}(x_i, x_j)$ becomes

$$E_{ij}(x_i, x_j) = \mu \hat{w}_{i,j} \quad (99)$$

where $\hat{w}_{i,j}$ is a weighting factor that assigns a penalty to all inconsistently labelled pair of neighbouring pixels (in order to produce a smooth curve). This energy was minimised over sufficient iterations to give a stationary solution. The mask function was then updated by generating a new shape from the segmentation. The energy was then minimised again for the new shape, the shape term was updated a fixed number of times to generate the best results.

The relative strength of the shape term compared to the region term is controlled by a coefficient.




Figure 32: Example image of graph cut segmentation result overlaid on original image.

3.8.1 Improvements to Graph Cut Model

The initial results were promising. A significant reduction in speed was achieved however the quality of the results was slightly worse using graph cut when compared to level set based methods. Therefore it was decided to continue investigating this method and find alternative ways of implementing a shape term in graph cut segmentation. Two different methods were attempted in order to improve the results these were using a boundary term to represent the shape and using a distance function instead of the binary function used previously. The use of a boundary term produced poor results so this technique was not adopted for use. Using a distance function resulted in an improvement in the segmentation.

3.8.2 Using Boundary Term to Generate Shape

Freedman and Zhang published a paper that produced an interactive graph cut model which used shape information to improve segmentation [128]. Their model was an interactive model, while I am trying to produce a model that involves no user

input. There are, however, interesting parts of their model. They attempt to introduce shape information by modifying the energy function

$$E = (1 - \lambda)E_i + \lambda E_s \quad (100)$$

where E_i is the image energy and E_s is an energy term based on a shape prior. It is the construction of the shape energy term E_s that is of particular interest. In order to set up the shape equation they set up a distance function $\bar{\Phi}$ with a zero level set corresponding to the shape template. Where the shape template is \bar{c} the following holds true

$$\bar{c} = \{x \in R^2: \bar{\Phi}(x) = 0\} \quad (101)$$

It is worth noting that $\bar{\Phi}$ is an unsigned distance function in contrast to the signed distance functions used earlier.

The shape energy is then given as

$$E_s = \sum_{(i,j) \in N: A_i \neq A_j} \bar{\Phi}\left(\frac{i+j}{2}\right) \quad (102)$$

where N is neighbouring set of pixels. The energy of this function will be low if $\bar{\Phi}\left(\frac{i+j}{2}\right) \approx 0$. If a point x lies near the template shape then $\bar{\Phi}(x) \approx 0$ and if the point lies far away from the boundary of the segmented object then $\bar{\Phi}\left(\frac{i+j}{2}\right) \approx 0$. Hence this energy is a function that encourages the boundary of the evolving contour to lie near the boundary of the shape function.

Using this method of creating the shape function, the function $E_i(x_i)$ now has no dependency on the shape function and becomes

$$\begin{aligned} E_i(x_i) &= |I(p) - c_1| && \text{inside the contour} \\ E_i(x_i) &= |I(p) - c_2| && \text{outside the contour} \end{aligned} \quad (103)$$

The shape function is instead found in the $E_{ij}(x_i, x_j)$ which becomes

$$E_{ij}(x_i, x_j) = \mu \hat{w}_{i,j} + \lambda \bar{\Phi}\left(\frac{i+j}{2}\right) \quad (104)$$

This model was implemented using my shape term, but the results were not as good as those produced using the binary shape graph cut method described previously.

3.8.3 Using a Distance Function for Shape Term

In the level set segmentation technique a signed distance function was used to represent the shape term. Initially this was replaced by a binary shape term for the graph cut segmentation work. The results using this produced less accurate results than the level set technique. In order to improve this technique it was decided to try using a distance function in graph cut work. An unsigned distance function was used here rather than a signed distance function since negative costs are not allowed in graph cut methods.

The initialisation steps for this technique remained the same as those used before. The difference was instead of generating a binary mask from the initial estimate of the shape an unsigned distance function was instead produced. The general energy function remains the same as

$$E(x) = \sum_i E_i(x_i) + \sum_{i,j \in N} E_{ij}(x_i, x_j) \quad (105)$$

However equation 100 for $E_i(x_i)$ now becomes

$$E_i(x_i) = |I(p) - c_1| - \lambda_{shape} * \left(\frac{1}{I(p)}\right) * D(p) \quad \textit{inside the contour} \quad (106)$$

$$E_i(x_i) = |I(p) - c_2| + \lambda_{shape} * \left(\frac{1}{I(p)}\right) * D(p) \quad \textit{outside the contour}$$

where c_1 and c_2 are mean intensities for inside and outside contour, $I(p)$ is image intensity at p , λ_{shape} is the shape weighting term at p and $D(p)$ is the value of the shape distance function at p . The $E_{ij}(x_i, x_j)$ term is unchanged compared to the binary shape version. The minimisation of this function was carried out the same way as before.

Results of this method are discussed in chapter 5.

3.9 Graph Theory Segmentation

The next approach to find a fast accurate segmentation method was by investigating approaches using graph theory. As discussed in Chapter 2 this is a technique that has been previously used to successfully segment both corneal and retinal OCT images.

Work using graph theory to segment images can be split into two different categories based on how the energy or cost function is minimised. In the first category are those approaches using edge weights, and minimising the energy using either Dijkstra's algorithm or Nilsson's A algorithm (two closely related methods). The other category of methods is based on using dynamic programming to find the shortest path. This can be optimised using either edge or node weights.

The pre-processing techniques used for this method differed slightly to the previous methods. This method is more robust to incorrect rotation of the images so the step to align the images horizontally was not needed. The removal of the central noise was also not carried out as this did not cause problems for the segmentation.

3.9.1 Dijkstra Algorithm Method

The first attempt at graph theory segmentation was based on using edge weights and Dijkstra's method to find the shortest path. This is similar to the work by LaRocca et al. [97]. For this application a novel shape term was added to improve the segmentation. The results of segmentation with and without the shape term were compared.

A combination of using the gradient of the image and a shape term was used to create the cost function. The edge going from node a to node b was given the weighting

$$w_{ab} = 2 - (g_a + g_b) + \lambda_s \Phi + w_{min} \quad (107)$$

where w_{ab} is weight of edge connecting a to b , g_a is vertical gradient at node a , g_b is vertical gradient at node b , Φ is a distance function that defines the desired shape, λ_s is a coefficient that determines the strength of the shape term and w_{min} is a small minimum value added for stability.

The shape function that is chosen is generated during an initialisation step. The image is segmented using a threshold, with the value of the threshold being determined using Otsu's method. While this method gives a good result for the location of the anterior surface of the cornea, it is not able to determine the lower surface. A 4th order polynomial is then fitted to the anterior surface. This polynomial is shifted down to give an estimate of the location of the lower surface. The distance it is shifted is related to the central cornea thickness, which can be detected from the image. An additional distance related to the distance from the centre is also added to account for the fact the cornea is thinnest at the centre. These two boundaries are then used to generate a distance function that has a

minimum on the boundary, and so acts to constrain the line produced to be close to the initial estimate of the location of the cornea.

After the weightings of the graph nodes have been determined, it is required to set end points for the image. Due to the curved nature of the cornea and the tendency of graph theory based methods to find straight lines as a shortest path it was desirable to set the start and end points to be fixed to the bottom left and bottom right of the image. This was done by fixing the first three pixels in the corner of each image to have a very low weight, making them attractive for the segmentation to detect.

In addition to this step, it was necessary to remove a section of the middle of the image from the image. This was to prevent the method finding a solution at the bottom row of the image. The section removed was at a fixed distance below the estimate of the anterior surface. A distance of three times the estimated corneal width was chosen for this since this is sufficiently large to definitely not contain any part of the cornea.

Once the weighting term and the endpoints were set, the shortest path was found using Dijkstra's algorithm. This was done separately for positive and negative gradients because the anterior boundary is defined as a shift from dark to light while the posterior boundary has a shift from light to dark.

3.9.2 Alternative Approach using Dynamic Programming

Although the results of the segmentation using the above technique were good, the speed of the program is still an issue. Whilst the above program is quicker than the level set method it is still not as fast as the graph cut method, discussed in the section 3.5. It is therefore desirable to attempt to speed up this technique.

Dynamic programming was investigated as an alternative method of minimising the energy function. A move to a node based cost system was also made since this makes calculating the weighting function simpler and less memory intensive.

A five stage segmentation method was developed. Figure 33 shows the different steps involved in the process. These will be now discussed in order.

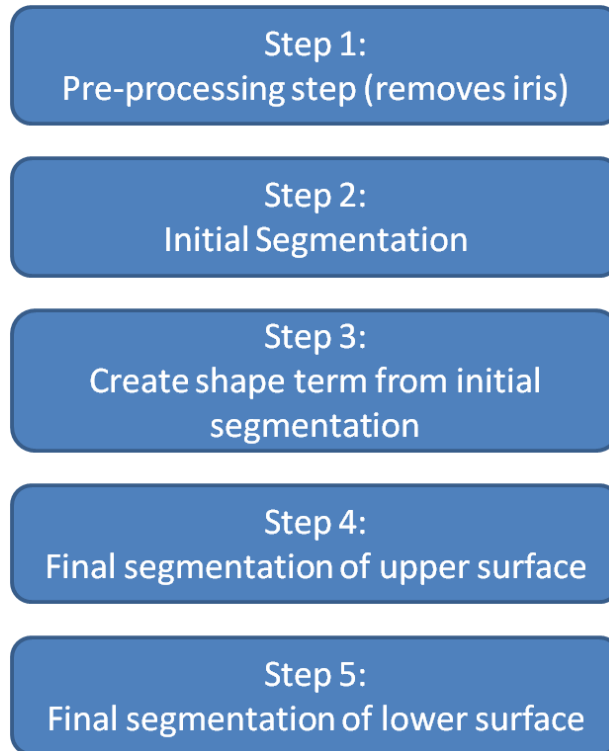


Figure 33: Flow chart showing the steps of the graph theory segmentation method developed.

3.9.2.1 Pre-processing Step

The first stage of the segmentation was the detection and isolation of the iris. The detection method used was the same as that described earlier, in section 3.5.1.

Once the iris is detected, this area is excluded from the next stage of the segmentation. The reason for removing the iris is that, as it is a horizontal structure in the image it would be detected instead of the cornea for some images.

The removed iris was allowed to be at different angles. Figure 34 shows the image after this step has been carried out.

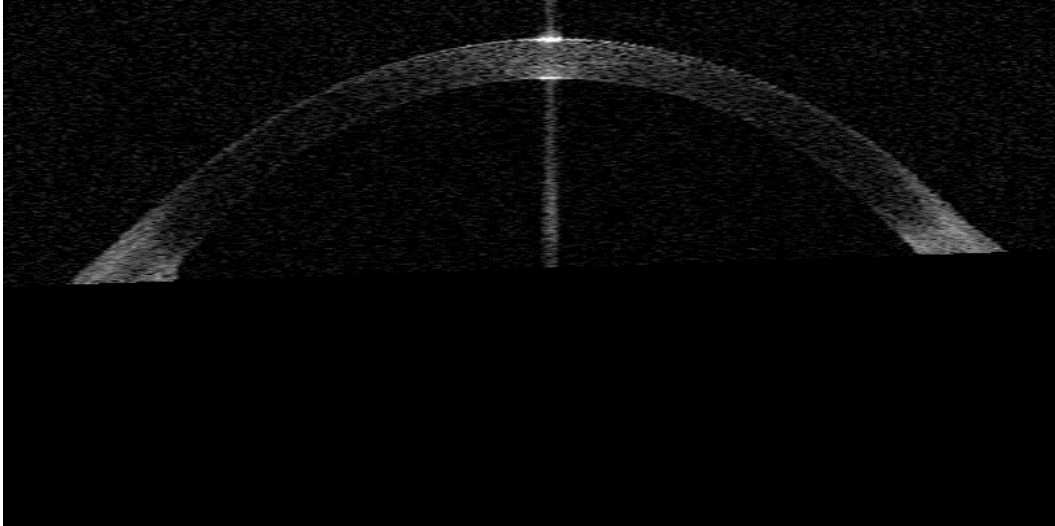


Figure 34: Image after the section containing the iris has been removed. The intensity of the removed black region was set to infinity to prevent any pixels in this area being included on path.

3.9.2.2 Initial Segmentation

In order to calculate the shortest path using this method, a cost function must be created that determines the cost of each pixel. A different function was used for each step of the segmentation. For the initial finding of the anterior surface the following cost function was used

$$C_1(i, j) = w_1 * Gradv(i, j) + w_2 * Gradh(i, j) \quad (108)$$

where $w_{1,2}$ are weighting coefficients, $Gradv(i, j)$ is the vertical gradient of the image at point (i, j) and $Gradh(i, j)$ is the horizontal gradient of the image. For the horizontal gradient the central region of the image was set to zero. The reason for using a combination of the two gradients is the curvature of the cornea. In the central region the vertical gradient is at correct orientation to find the boundary. As the cornea curves towards the edges of the image the boundary is orientated towards the top corners of the image so using both horizontal and vertical gradient is important.

When carrying out this initial segmentation the image with the iris removed was used. The results of this segmentation is shown in Figure 35 plotted on the original image.




Figure 35: Image with result of initial segmentation step marked on it. Note the segmentation was carried out on the cropped image but is plotted on the full image here. This is why the line doesn't follow the obvious boundary to the edges of the image.

3.9.2.3 *Creation of shape term from initial segmentation*

Excluding the iris also removes the section of the cornea where it joins the sclera. It is therefore desirable to segment the entire image without removing the iris. Segmentation of the posterior surface is also necessary and as previously discussed suffers from regions of low signal to noise ratio that cannot easily be detected.

In order to aid further segmentation, two shape functions were created. For the detection of the anterior surface a function ϕ_1 is created. This is a distance function from the anterior boundary detected in part 2 of the segmentation. The value for the function at any point is simply the distance of that point to the line. This function encourages subsequent segmentations to follow the path of the initial segmentation.

For the lower surface, the upper surface of the cornea was used as a guide. A model shape term was then used to estimate the width of the cornea. In order to do this the width of the cornea at the centre was first measured. The width of the cornea was found by identifying peaks in the intensity of slices of the image close to the centre. Only peaks with values of 50% of the largest peak were considered. This threshold was used as due to noise there are many local maximum, local maximum that are due to noise rather than image information are much smaller so are removed by this step. The anterior surface of the cornea was assumed to be the first significant peak that was found. The lower surface of the cornea was

assumed to be the last significant peak. Any peaks that were too close to the bottom of the image were excluded since they can't be part of the cornea. This technique only works in the central region where the best signal to noise ratio is found.

Once the central cornea width was detected, a 2nd order polynomial was used to model how the width of the cornea increases away from the centre. For simplicity of calculation the vertical width rather than actual width of cornea was used. The coefficients of the polynomial were determined empirically by testing different values and seeing what produced the best results.

A distance function was then created from this estimate of the location of the posterior surface of the cornea. The motivation for using this shape function rather than fitting an ellipse as done previous was to improve the accuracy of the method. The ellipse is not a perfect model for the cornea. The graph theory method by it's nature gives us an accurate segmentation of the anterior surface which could be used to give a more accurate estimation of the position of the posterior surface than an ellipse.

3.9.2.4 *Final Segmentation of the Anterior Surface*

Once a shape term had been generated from the initial segmentation, a new cost function was constructed for the anterior surface and then the segmentation was carried out again. The new cost function is given by

$$C_2(i, j) = w_1 * Gradv(i, j) + w_2 * Gradh(i, j) + w_3 * \Phi_1(i, j) \quad (109)$$

where $\Phi_1(i, j)$ is distance function that determines the distance from the line detected in the segmentation of the image with the iris removed. This encourages the line to follow the path of the previous segmentation but allows it to extend into regions that were not part of the previous segmentation. An example of the result of this step is shown in Figure 36.




Figure 36: Image with the result of the final segmentation of the anterior surface using graph theory marked on it.

3.9.2.5 *Final Segmentation of the Posterior Surface*

The next step of the segmentation is to detect the posterior surface. This is harder to detect than the anterior surface as the boundary is not as well defined. In order to aid the detection of the boundary two strategies were used: the image was flattened and a shape term was incorporated in the energy function. Flattening the image aids accurate segmentation since the graph theory segmentation technique favours detection of straight lines. The image was flattened using the results of the anterior surface segmentation. Each column of the image was moved so that the anterior surface detected in the previous step is a flat line in the image. Figure 37 shows what an example image looks like after it has been flattened.




Figure 37: Example image that has been flattened using the segmentation of the anterior surface shown in Figure 36. Since the shift applied to each point is known it can be very easily reversed to align the results with original image afterwards.

The detection of the lower surface was carried out in a similar manner. An energy function incorporating a shape weighting term was used. The weighting function was given by

$$C_3(i, j) = w_1 * (-1) * Gradvfilt(i, j) + w_2 * \Phi_2(i, j) \quad (110)$$

where $w_{1,2}$ are weighting functions to determine the strength of the different terms, $Gradvfilt(i, j)$ is the gradient in the vertical direction after a median filter was applied to the image the inverse of gradient was used since the posterior boundary is a light to dark edge. This filter was used to reduce to effect of noise on the segmentation. $\Phi_2(i, j)$ is the shape term determined from the anterior surface through the method described above.

In addition to this regions that were definitely known to not be part of the segmentation were given infinitely high values to remove them from the possible segmentation. This was determined by being more than half the corneal width above the expected value, which prevents a second detection of the anterior corneal surface. An area more than the corneal width below the expected value was also excluded to prevent detect of the iris instead of cornea. Figure 38 shows the results of the segmentation of anterior and posterior boundaries.




Figure 38: Image with anterior and posterior boundary marked on it.

3.9.2.6 Energy Minimisation

Following the determining of the weighting function the shortest path was found using dynamic programming technique. This works by calculating the cumulative cost function for each point using the following expression

$$t(i, j) = \begin{cases} \infty & j < 1, j > m \\ C_q(i, j) & i = 1 \\ \min_{p=j-3:j+3} \begin{pmatrix} t(i-1, j-3) + C_q(i-1, j-1) + C_q(i, j), \\ t(i-1, p) + C_q(i, j), \\ t(i-1, j+3) + C_q(i-1, j+1) + C_q(i, j) \end{pmatrix} & \text{otherwise} \end{cases} \quad (111)$$

where $t(i, j)$ is the cost to reach point (i, j) , $C_q(i, j)$ is the cost accumulated at node (i, j) the q designates which cost function is used this minimisation scheme is used for all three cost functions described above, i is the x direction index and j in y direction index with maximum value m . This allows each pixel to be connected to one of 7 pixels from the previous row. If it is one of the 5 closest only the values at the previous point is taken into account. If there has been a move of 3 pixels up or down then an additional cost is added from a point that it is assumed the line has to pass through in order to make a move that distance. This is done to allow for steep lines to still be detected while preventing too much movement in the vertical direction without a cost penalty occurring. Only one point is found for each point on the x axis. Since the aim of the segmentation is to find a structure that spans the image this is a desirable condition for the path.

One of the advantages of this system is that it doesn't require the use of a start or end point. The shortest path that connects any pixel where $i = 0$ to any pixel with $i = n$ is found.

Once both boundaries were segmented they were both smoothed using moving average filter. This is to remove small amounts of noise in result due small variations in pixel intensity along boundary. The results on this can be seen in Figure 39 below.




Figure 39: Image with final segmentation of anterior and posterior boundaries. Here the boundaries have been smoothed and the posterior boundary cropped to stop at the boundary with the iris.

3.9.3 Detection of Additional Features

How to detect the anterior and posterior surfaces of the anterior segment using this technique has been described. Within the anterior segment there are also a number of other points of interest. These points of interest are related to the iris. The addition of the segmentation of the iris is therefore desirable. This can be done using a variation of the technique used to find the other boundaries.

The pre-processing step found a straight line estimate of the location of the iris based on finding a line of maximum brightness in an image. This was useful when removing the iris but in order to segment it, a more accurate method is needed. To find the top surface of the iris a slight variation to the energy term used to find the anterior surface of the cornea can be used. The weighting function used to detect the top surface of the iris is given by

$$C(i, j) = w_1 * Gradv(i, j) + w_2 * \Phi_{iris}(i, j) \quad (112)$$

where $Gradv(i, j)$ is the vertical gradient of image, $\Phi_{iris}(i, j)$ is a shape term generated from the straight line estimate of the iris and w_1, w_2 are coefficients governing the strength of the different terms. The guide shape term is zero in a region close to the approximate location of the iris as found in the pre-processing step and valued at infinity everywhere else. This acts to prevent any segmentation that includes points that are not close to initial estimate of iris location. This was done to prevent detection of the cornea a second time. The iris is the strongest structure on the image other than the cornea.

Finding the lower boundary of the iris is more difficult. This boundary is not well defined in the OCT images and lacks a strong edge that has been used to detect the other boundaries. In order to detect this, an additional term is added that looks at the intensity of the points below this in the image. The bottom boundary of the iris should have an area of darkness below. The weighting function used is given by

$$C(i, j) = w_1 * Gradminusv(i, j) + w_2 * \Phi_{iris}(i, j) + w_3 * darkbelow(i, j) \quad (113)$$

Where $Gradminusv(i, j)$ is the negative gradient as used before, $\Phi_{iris}(i, j)$ is same shape term used for top surface and $darkbelow(i, j)$ is new term that is a cumulative total of pixel intensity for each row favouring points lower on the image and those below any bright sections of the image. Each term has a weighting factor w_k that governs its strength.

Using these two weighting functions it is possible to obtain boundaries for the iris. The method always segments continuous lines from left to right in an image. This means that an additional section will be segmented joining the two sections of the iris together. This boundary in between the two sections of the iris does not show anything useful. In some images, sections of the lens are visible and in these cases the segmentation will follow this but it is not reliable enough to be used to detect the lens. The lens is not visible in all images.

There are two keys points that it is useful to identify: the junction of the iris and the cornea (ie. the apex of the anterior chamber angle) and the edges of the pupil. The junction of the iris and cornea can be located by finding the intersection of the posterior boundary of the cornea with the anterior surface of the iris. The edges of the pupil are more difficult to find. This was done by looking at the gradient in the x direction near the top surface of the iris. The end of the iris should correspond to

the maxima of the absolute value of this gradient. The image was split into two halves and the maxima closest to the centre on each side were taken as the two ends of the iris. The pupil is the region between the two sections of the iris so finding the end of the iris and the edge of the pupil are the same thing.

Figure 40 is an example image with the segmentation of cornea and iris marked on it. In the middle region of iris segmentation nothing of interested is segmented.




Figure 40: Example image with segmentation marked on. Red lines locate the anterior and posterior boundaries of the cornea. Green lines are segmentation of iris. The red asterisks mark the located points of interest on this image. Both the apex of anterior chamber angles are detected on both sides and the two edges of the pupil are also detected in the middle.

3.9.3.1 Improvements to Iris Detection

There were two main problems faced when carrying out segmentation of the iris using graph theory. The first is the discontinuous nature of the iris. The graph theory technique can be used best to segment complete structures that connect one side of an image to the other. This means that when investigating the iris there is a need for an additional step of the detection of the ends of the iris. The other problem is the weak edges, especially on the lower edge of the iris. The technique described above uses edge information to find the boundary, so the lower boundary of the iris is liable to achieve incorrect segmentation.

In order to improve the technique the previous iris segmentation technique, discussed in section 3.7.6 was used. It was altered slightly since a segmentation of the posterior boundary of the entire image has already been achieved. First the region above the posterior boundary of the cornea was removed from the image in order to prevent any repeat detection of the cornea. The remaining image was then segmented using a threshold set by Otsu's method. A morphological closing operation was then carried out to fill in small gaps in the segmentation due to noise.

The two largest objects detected are either side of the iris. This can be combined with the previous segmentation to include the iris in our segmentation. The junction between the iris and the cornea can be found by finding the minimum horizontal distance between the points found on the iris and the posterior segment. Where this is a minimum, the junction between the cornea and the iris has been found. Figure 41 shows an example image with segmentation of boundaries of cornea and iris marked on.

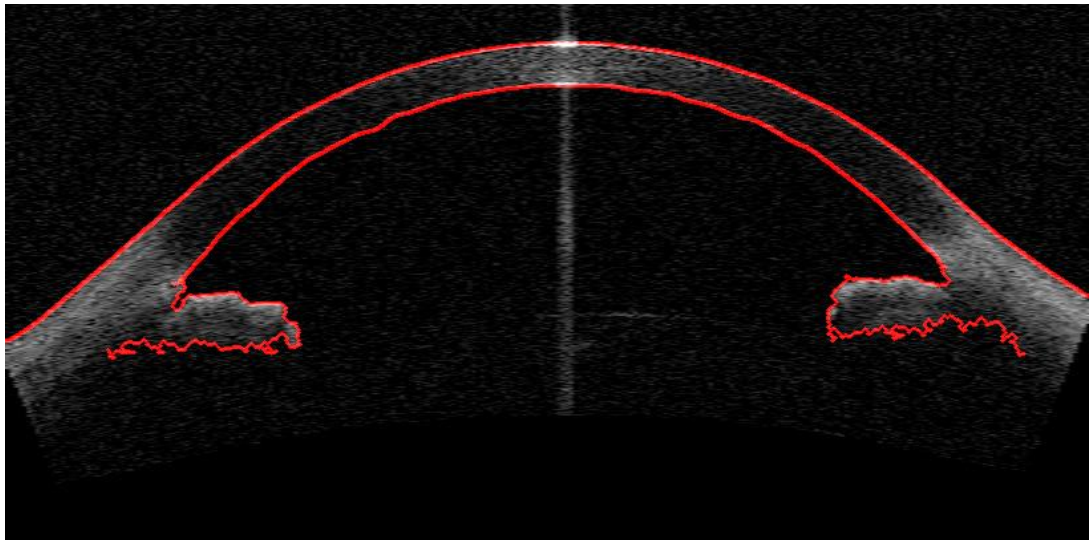


Figure 41: Image with segmentation of iris and cornea marked on.

3.9.4 Using C++ to Improve Speed of Segmentation

Initially the graph theory segmentation program was implemented in Matlab. Matlab was chosen for the ease of use. Following positive evaluation of the graph theory with dynamic programming technique it was decided to additionally implement the technique in C++. This was done in order to increase the speed of the program. The technique was unchanged when it was re-coded from Matlab to C++. These two versions of the function were validated separately and the results compared.

3.10 Post Processing

After segmentation of the anterior and posterior boundaries had been carried out the boundaries produced were often not smooth. This is due to speckle noise inherent in OCT imaging and the discrete nature of the boundary produced. In order to produce a boundary that is physically realistic smoothing must occur.

Initially this was carried out by fitting a polynomial to the surfaces generated. A 10th order polynomial was fitted to the anterior surface and a 4th order polynomial was fitted to the posterior surface. These polynomials were chosen since they have both been previously used to smooth corneal surfaces.

Another smoothing technique that was also used was a moving average filter. This works by making each point on the curve the mean value of its neighbours.

3.11 Conclusion

A number of different segmentation techniques have been developed. These can be split into three main categories; level set techniques, graph cut segmentation, and graph theory segmentation. A number of variations of each method have been developed.

The introduction of a shape term was found to be important for all the different segmentation techniques. This shape term allowed the method to overcome problems due to parts of the image having a low signal to noise ratio. The implementation of a shape term for segmentation of anterior segment OCT images is the main unique contribution of this part of the study.

Two different methods have been used for creating a shape term. Initially an ellipse was fitted to the surface to provide a guide shape for the cornea, this shape was used for the level set and graph cut techniques. An ellipse was chosen because the structure of the eye can be described as two ellipsoids. The main section of the eyeball is one ellipsoid and then another can be used to represent the cornea which extrudes from this other ellipsoid. When using an iterative approach such as level set or graph cut the shape term was updated repeatedly throughout the segmentation to generate new ellipses as the segmentation proceeds. The main motivation for using a different shape term for the graph theory segmentation technique was limitations of using an ellipse to describe the cornea. The approximation of the cornea being an ellipse is only valid in the central section at the limbus the cornea merges into the sclera. The sclera has a different curvature than the cornea so the same ellipse can be used to model the surface of this. The anterior segment images being used did include small sections of the sclera so the elliptical shape term became less accurate at the edges of the images.

The step nature of the graph theory technique where each surface is found in turn motivated a different approach. Having a good segmentation of the anterior surface prior to segmentation of the posterior surface, which was much easier to find and didn't require the use of a shape term, meant that this could be used to generate a more accurate shape term instead. The central thickness of the cornea could be calculated on the image utilising the much better signal to noise ratio in the central part of the image. This could then be used to estimate the location of the posterior

surface on the rest of the image which was used as a shape term to guide segmentation. This method proved more accurate than the use of the ellipse.

It would have been possible to use the initial segmentation that was used to initialise the level set technique to produce a similar shape term for this technique. This was not done since the graph theory segmentation technique was preferred for reasons of speed and memory usage.

In order to evaluate the methods the segmentation of the same set of 39 images was carried out using all the different techniques. These images have also been manually segmented by two expert ophthalmologists. The results of this evaluation are described in Chapter 5.

4 3D Segmentation

4.1 Introduction

Chapter 3 discussed a number of 2D segmentation methods. However, the cornea is a 3D structure and in order to produce a reliable model of the cornea 3D segmentation must be carried out. To do this a suitable imaging framework is needed that is capable of producing a 3D segmentation. Registration of multiple 2D images of the same object must also be considered.

In this chapter the expansion of each of the 2D segmentation techniques developed into 3D is discussed. First an attempt at expanding the level set with shape segmentation technique to 3D is discussed. This method was found to require too much memory and be too slow to be a practical segmentation technique, which motivated the investigation of other segmentation techniques. Following this the expansion of the graph cut segmentation method to 3D is carried out. Finally the expansion of graph theory segmentation to 3D is discussed. Registration of images is discussed as a part of the techniques.

4.2 3D Level Set Segmentation

A 2D segmentation method using a level set function has been developed. One of the advantages of using a level set method is the ease of extension to include 3D data sets. There are a few modifications needed to make the technique applicable to 3D images but the basic model used is very similar. The three step algorithm used for 2D segmentation is used with slightly different pre-processing steps since images are different.

3D AS-OCT images using a customised AS-OCT device were used for this technique. The dataset comprises 32 radial images centred at the cornea with an angle of 11.25 degrees between images. For this specific problem using cylindrical polar coordinates is more appropriate than Cartesian coordinates due to the method by which the images are acquired. Level set models can be extended into 3D without changing the energy function fundamentally. Previous work on 3D segmentation has mainly focused on models using Cartesian coordinates (i.e. a series of parallel scans of a subject being used to create a 3D image) [101].

4.2.1 Segmentation Framework

A three step segmentation framework is developed. The first step is to pre-process the image to reduce noise. This involves using a combination of Gaussian filters and morphological processing to smooth noise from the image. The next step is to

obtain a coarse segmentation of the anterior segment. The final step uses the new level set based shape prior segmentation model to evolve the contour initialised from the coarse segmentation and achieve the final segmentation.

4.2.1.1 **Pre-processing Step**

The initial pre-processing step is to apply a median filter to the image. This reduces the noise in the image and is a filter that works well to remove the speckle noise found in OCT images. Morphological processing is then used to remove some of the unwanted structures in the image. There is a tendency for bright horizontal bands to form above the cornea in the image, as can be seen in Figure 4b. These are removed by a morphological closing operation. Linear structural elements are used in this process.

4.2.1.2 **Coarse Segmentation**

The aim of this step is to produce an initial estimate of the corneal location (or coarse segmentation). This estimate is used as the initial location of the curve to be evolved by the level set function in the following step. Additionally the anterior boundary is used to construct the shape constraint in the later stage. A similar approach can be taken to the one used previously in 2D segmentation. More specifically, an entropy filter is applied to the pre-processed image to produce an entropy map. The coarse segmentation is achieved by segmenting the entropy map using Otsu's thresholding method. This is applied to the first image and a full initial shape is created by rotating this image, creating a perfectly cylindrically symmetric initial shape.

4.2.1.3 **Segmentation with Level Set and Shape Prior**

A general segmentation model using level set and shape prior can be described as the following energy minimisation problem

$$\begin{aligned}
 E(\Phi(\mathbf{x}), c_1, c_2) = & \lambda_1 \int_{\Omega} \delta(\Phi(\mathbf{x})) |\nabla \Phi(\mathbf{x})| d\mathbf{x} \\
 & + \lambda_2^1 \int_{\Omega} (I(\mathbf{x}) - c_1)^2 H(\Phi) d\mathbf{x} + \lambda_2^2 \int_{\Omega} (I(\mathbf{x}) - c_2)^2 (1 - H(\Phi)) d\mathbf{x} \\
 & + \lambda_3 \int_{\Omega} (\Phi(\mathbf{x}) - S(\mathbf{x}))^2 d\mathbf{x} + \lambda_4 \int_{\Omega} (|\nabla \Phi| - 1)^2 d\mathbf{x}
 \end{aligned} \tag{114}$$

where $\Phi(\mathbf{x})$ is the level set function, $\delta(\Phi)$ and $H(\Phi)$ are the delta function and Heaviside function respectively, $I(\mathbf{x})$ the image intensity, c_1 and c_2 the mean intensities inside and outside the zero level contour, and Ω the entire image volume. λ s are the weighting parameters for different terms. In particular λ_2^1 and λ_2^2 can be

used to apply different weights to the two region terms. $S(\mathbf{x})$ is a level set function corresponding to the shape prior. The first three terms in equation 114 are the standard Chan and Vese model [70]. The fourth term is a shape term that keeps $\Phi(\mathbf{x})$ close to the shape prior $S(\mathbf{x})$, the formulation of $S(\mathbf{x})$ will be detailed later. The last term is a regularisation form introduced by Li et al [80] to keep $\Phi(\mathbf{x})$ as a valid level set function without need of computationally expensive re-initialisation. The effect of this term is to penalise the formation of any regions with either very steep gradient or any flat areas. The formation of these regions would mean the function was no longer a valid level set and so require reinitialising.

The regularisation term is introduced to this model as, when carrying out 3D segmentation reducing time taken for segmentation became more important. This is due to the large size of the images giving high memory requirements for the model. By including this term the computationally costly re-initialisation step can be removed from the segmentation.

The solution to equation 114 can be derived by a gradient descent approach as follows

$$\begin{aligned} \frac{\partial \Phi}{\partial t} = & \lambda_1 \operatorname{div} \left(\frac{\nabla \Phi}{|\nabla \Phi|} \right) \delta(\Phi) - \lambda_2^1 \delta(\Phi) (Z(\mathbf{x}) - c_1)^2 + \lambda_2^2 (Z(\mathbf{x}) - c_2)^2 \\ & - 2\lambda_3 (\Phi(\mathbf{x}) - S(\mathbf{x})) + \lambda_4 (\nabla^2 \Phi - \operatorname{div} \left(\frac{\nabla \Phi}{|\nabla \Phi|} \right)) \end{aligned} \quad (115)$$

The energy function was minimised by alternatively minimising Φ , c_1 and c_2 . When Φ is fixed, the terms c_1 and c_2 are found using the following expressions

$$c_1 = \frac{\int_{\Omega} Z(\mathbf{x}) H(\Phi(\mathbf{x})) d\mathbf{x}}{\int_{\Omega} H(\Phi(\mathbf{x})) d\mathbf{x}} \quad (116)$$

$$c_2 = \frac{\int_{\Omega} Z(\mathbf{x}) (1 - H(\Phi(\mathbf{x}))) d\mathbf{x}}{\int_{\Omega} (1 - H(\Phi(\mathbf{x}))) d\mathbf{x}} \quad (117)$$

where $Z(\mathbf{x})$ is the image intensity at point \mathbf{x} and $H(\Phi(\mathbf{x}))$ is the Hausdorff function which determines is a point is inside or outside of the boundary.

The shape prior, $S(\mathbf{x})$, was updated once every 100 iterations using the following process. An ellipsoid was estimated by least square fitting to the anterior surface from the level set function. A second related ellipsoid was built at a fixed distance

below the first one. The shape prior of the cornea $S(\mathbf{x})$ was then computed as the product of the level set functions corresponding to those two ellipsoids.

$$S(\mathbf{x}) = S_{upper}(\mathbf{x}) S_{lower}(\mathbf{x}) \quad (118)$$

where $S_{upper}(\mathbf{x})$ is a signed distance function corresponding to ellipsoid fitted to anterior surface and $S_{lower}(\mathbf{x})$ corresponds to a related ellipsoid which has been shifted down to mimic the lower surface of the cornea. Taking the product ensures a sign difference between the volume between the ellipsoids and outside the ellipsoids.

In the discretisation, the main difference between cylindrical and Cartesian coordinate is the curvature term: the former one is more complex. Specifically, the curvature under cylindrically polar coordinate becomes

$$div\left(\frac{\nabla\Phi}{|\nabla\Phi|}\right) = \frac{1}{r} \frac{\partial}{\partial r} \left(\frac{r \Phi_r}{|\nabla\Phi|} \right) + \frac{1}{r} \frac{\partial}{\partial \theta} \left(\frac{1}{r} \Phi_\theta \frac{1}{|\nabla\Phi|} \right) + \frac{\partial}{\partial z} \left(\Phi_z \frac{1}{|\nabla\Phi|} \right) \quad (119)$$

where r , θ and z are cylindrical coordinates. This was implemented using central difference approximations for the partial differentials in 3D.

The level set function $\Phi(\mathbf{x})$ was updated until either the update has only a very small effect on the position of the contour or 2,000 iterations has been carried out.

Major problems faced with this method are those of speed and memory usage. When attempting to implement this method using level set technique using images sized 32x2048x2048 pixels it was found the computer lacked sufficient memory to run the program. In order for segmentation to proceed the size of the images was reduced. Finding alternative approaches that can speed up the segmentation and reduce memory usage is therefore necessary.

Another limitation this method suffers from is that it fails to consider potential problems with the registration of the images. This was not an issue when synthetic data was used however when real data was used it became apparent this was a problem. The problems with slow speed, lack of registration of images and the need to make images smaller for segmentation to succeed meant that work on this method was not pursued further. No attempt was made to carry out the evaluation of this technique using the repeated data set since the technique had been ruled out on grounds of speed and memory usage prior to the acquisition of the repeated data set.

4.3 3D Graph Cut Segmentation

Graph cut techniques have been shown to quickly segment 2D AS OCT images effectively. The technique used can be expanded to segment 3D images. The main difference when moving from 2D segmentation to 3D segmentation is the need for the addition of a registration step before segmentation can take place. The longer time taken for imaging in 3D means that there is an increased chance of movement between the scans. This means that an additional preliminary step to align the images before segmentation needs to be carried out.

4.3.1 Pre Segmentation Alignment

In order to register the images they are first segmented individually on a 2D basis. This is done using the 2D graph theory method described in section 3.9.2, applied to the anterior surface only. This approach is used since it is a quick, easy and accurate method. The time taken was reduced by only finding the location of the anterior surface. Once the anterior surfaces of the set of 16 images were acquired they were aligned to each other by finding the top point on each surface and aligning these.

Initially the images are cropped to only include the radius of the cornea that is present in all the images. This is done as smaller images result in faster segmentation. While some information has been discarded this way, this is from areas that have a very low density of points so information there would be less reliable than in more central regions where there is a higher density of points.

A variation of the method is to alter where the image were cropped off. There was a requirement for all images to be the same size in order to combine them to a 3D volume. This means that cropping less off some images requires adding zero padding onto other images. Using larger images will results in an increase in time taken for segmentation but this allows for segmentation to be carried out over a larger area. Two different distances were used to see what effect this process has on segmentation.

It would be possible to carry out further tests to develop better registration techniques for this method. This was not done as the analysis of 2D images found that graph theory segmentation was more accurate than graph cut segmentation. 3D segmentation efforts were therefore mainly focused on this method.

4.3.2 Creation of Shape Term

Once the images have been aligned a shape prior term needs to be created. This is done using the graph theory segmentation of the anterior surface which had already carried out for the purpose of alignment. An estimate of the posterior surface is found from the line. This is done in the same way the shape term is generated for the lower surface of the graph cut model (see section 3.9.2.3). The distance at the centre is calculated from intensity peaks in the image then the distance was assumed to increase either side. A signed distance function is created from the two lines. The area inside the two lines is given negative values, and the outside the lines, a positive value. This is done on a 2D basis separately for each image. These are combined to produce a 3D shape term.

4.3.3 Segmentation

The switch to a 3D graph cut segmentation alters the handling of the neighbourhood term as working in 3D affects the number of the neighbours each point will have when compared to working in 2D. The rest of the energy function is unaffected. With the creation of a suitable 3D shape term a similar shape term to the one used for 2D segmentation can be used.

The energy function therefore remains as

$$E(x) = \sum_i E_i(x_i) + \sum_{i,j \in N} E_{ij}(x_i, x_j) \quad (120)$$

where

$$E_i(x_i) = |I(p) - c_1| - \lambda_{shape} * \left(\frac{1}{I(p)}\right) * D(p) \quad \textit{inside the contour} \quad (121)$$

$$E_i(x_i) = |I(p) - c_2| + \lambda_{shape} * \left(\frac{1}{I(p)}\right) * D(p) \quad \textit{outside the contour}$$

$$E_{ij}(x_i, x_j) = \mu \hat{w}_{i,j} \quad (122)$$

where c_1 and c_2 are mean intensities for object and background respectively, $I(p)$ is image intensity at p , λ_{shape} is the shape weighting term at p and $D(p)$ is the value of the shape distance function at p .

Once the shape term is created and the 3D images are registered then 3D segmentation using graph cut is carried out. This is done using the same code that has been used for 2D segmentation. The only difference between the 2D and 3D cases is the neighbourhood term that is altered to reflect a larger neighbourhood for 3D images.

4.4 Graph Theory Segmentation in 3D

When investigating 2D images significant improvements in both speed and performance were achieved using graph theory segmentation. Graph theory methods cannot easily be expanded to 3D so instead a quasi 3D segmentation technique was used. The 2D images were segmented individually then a registration step was performed to align the images into a 3D model.

4.4.1 Segmentation

This is carried out in exactly the same as discussed in section 3.9.2 for 2D image segmentation with graph theory.

4.4.2 Registration

Following segmentation a registration step was performed. The relative angles of all the images were known as this is defined by the OCT system when the images were being taken. There is, however, a possibility of movement of the eye between scans causing misalignment between the images taken. The registration step needs to account for this effect. A number of different registration techniques are carried out to see which method performed best.

Initially relatively simple methods were attempted. Following this Genetic algorithm (GA) based techniques were developed. The GA based techniques proved most successful, the non GA based methods were used as an initialisation to the GA. This section will first discuss the GA registration technique. Other methods not using GA that were also attempted will then be explored.

4.4.3 Genetic Algorithms

Genetic Algorithms (GA) were chosen as a technique since they are able to easily find a global minimum of an energy function. They allow for complex energy functions to be for the purpose of registration. Initial attempts had used only a few points from the images to align the images. This has the advantage of producing a quick registration process; however more accurate registration may be achievable using more points from the images. In order to improve on these techniques a new

method was attempted that uses genetic algorithms (GA) to improve the registration.

GA work by breeding solutions from an initial set. A function is created that evaluates the goodness of a particular solution. Solutions that are a good fit are more likely to appear in the new batch of solutions than worse solutions and the best 4 solutions were always kept. The new solutions are then also evaluated and a further generation of solutions is generated. This process stops when no improvement occurs over a set number of generations. In this case the best solution was assumed to be reached if there was no improvement over 200 iterations. The method was also stopped after it had run for 2,000 iterations.

The Matlab implementation of the GA was used with a total population of 100, of which 4 are elite, 80 are crossover offspring and 16 are mutation offspring. These values are used since they are the default values in Matlab and altering them didn't produce a significant improvement in results.

One of the key tasks in the successful use of GA is the choice of a suitable energy function. A number of different energy functions were tested and evaluated against each other. In addition to the choice of energy function knowledge of what is a reasonable range of values to be investigated can also be used to improve this method. Reducing the solution space to only allow solutions in a particular area can act to improve the results

4.4.3.1 Construction of the Energy Function

In order to register the images it is necessary to exploit the fact information on the general shape of the object of interest is known. The aim is to obtain a shape that models the cornea. There are 16 slices arranged in a radial pattern and they need to be aligned to form a smooth surface. The energy function that is minimised should therefore give preference to producing smooth shapes that look like possible corneas.

A total of five different energy functions for genetic algorithms were compared. These were circles in r, circles in r and z, ellipse fitting, ellipse fitting with circles, and model fitting.

4.4.3.1.1 Circles in r

The first energy function used is based upon the idea that the cornea is roughly ellipsoidal in shape. An energy function based on the standard deviation of the height of points in concentric circles around the centre is used. This should act to

smooth any bumps caused by misalignment and result in a smooth cornea of roughly ellipsoidal shape. Five circles of different radius were used to minimise the effect of any anomalies in the segmentation results. The energy function is given by

$$E = \sum_{j=30:30:150} std(z|_{r=j}) \quad (123)$$

where $z|_{r=j}$ is the z values of points where $r = j$. j is varied at fixed intervals from points near the centre to further out. Figure 42 shows to locations of the points considered, the standard deviation the z values of points on each of the circles is calculated separately and the sum taken. Points that are a large distance from the centre are not considered since there is an increased chance of stray results further out. Each slice was allowed to move in the x , y and z directions. This gives a total of 48 coefficients to be optimised.




Figure 42: Image showing the energy function described here. The standard deviation of the heights of points located on each of the red circles is used as the energy function. GA is used to minimise this function and produce a surface with a series of circles of points with as close to uniform height as can be achieved.

The energy function is then minimised using genetic algorithms. The Matlab implementation of the GA is used with a total population of 100, of which 4 are elite, 80 are crossover offspring and 16 are mutation offspring. These values are used since they are the default values in Matlab and altering them didn't produce a significant improvement in results. The process was run until it showed no improvement over a period of 100 generations.

Results using this method are reasonable but the shapes produced contain noticeable ridges. Further investigations into ways to improve this energy function were made.

4.4.3.1.2 Circles in r and z

The initial method was based on fitting circles with fixed values of r and minimising the differences in z based on these. As an extension circles with fixed z values are also used and minimise the differences in r of these circles. This led to the addition of a second energy term

$$E_2 = \sum_{j=\min(z)+20:20:\min(z)+100} std(r|_{z=j}) \quad (124)$$

where $r|_{z=j}$ is the r values of points where $z=j$. The energy function used is a linear combination of this and the previously used function (equation 123). Giving

$$E = \lambda_1 \sum_{j=30:30:150} std(z|r=j) + \lambda_2 \sum_{j=\min(z)+20:20:\min(z)+100} std(r|_{z=j}) \quad (125)$$

where λ_1 and λ_2 are weighting functions to determine the strength of each term. They were kept equal.

4.4.3.1.3 Ellipse Fitting and Ellipse Fitting with Model

The third method used is to fit an ellipsoid to the data and minimise the distance between this ellipsoid and our surface. This has the effect of pushing the data towards being a smooth surface. An ellipsoid was used since this was found to be a reasonable approximation of the corneal surface during earlier segmentation work. The energy function for this became

$$E = \sum_{i=1:N} abs(z_i^{surface} - z_i^{ellipse}) \quad (126)$$

where $z_i^{surface}$ is a point on the surface, $z_i^{ellipse}$ is a point with same x and y coordinates on the ellipse and N is the total number of points. This new energy function was tried both on its own and in combination with the previously used energy function. This method has the disadvantage, when compared to previous models, of being substantially slower than the other methods. This is due to the fitting step needed to fit an ellipsoid to the surface. When using GA, the energy function must be repeatedly calculated, if using a population of 100 over 100 iterations then the energy function calculation is carried out 10,000 times. This means that it is desirable for the energy function to be able to be calculated very quickly.

4.4.3.1.4 Model fitting

Zernike polynomials have been extensively used to model the cornea. This makes them a likely candidate for use in the energy function. Instead of fitting an ellipsoid to the surface 5th order Zernike polynomials were used instead. A surface was

generated from the Zernike polynomial and the difference between the surface of the polynomial and the model points was used as the energy function. This gave a function of

$$E = \sum_{i=1:N} abs(z_i^{surface} - z_i^{zernike}) \quad (127)$$

where $z_i^{surface}$ is a point on the surface, $z_i^{zernike}$ is a point with same x and y coordinates on the polynomial surface and N is the total number of points. This energy function is used on its own. It was found that the method of fitting the Zernike polynomials was too slow for this method to be of practical use. The time taken for each step is 0.2s, with a population of 100 this meant it took around 20s for each iteration. Typically the code needs to run for around 600 iterations before it reaches the stopping criteria. The long time taken to run makes this method unsuitable for practical use: results from this technique were therefore not included in the analysis.

The use of a fixed surface as a model was considered next. Using the results from the Zernike method, which generated a smooth surface, a model surface is generated. This surface is then used instead of generating a surface from the points each step. This allowed for a much faster registration speed. The energy function stayed the same but the Zernike polynomial is not updated each step.

These five different GA energy functions were tested against each other using the same set of 17 eyes and a repeatability test.

4.4.3.2 Finding Centre of Rotation

The method of imaging the data meant that it consisted of a series of 4 sets of 4 images. Each image in the set was automatically rotated 45 degrees relative to the previous image around a common centre. Following this, a manual rotation of 12 degrees occurred and another set of 4 images was taken. This meant that each set of 4 images had a common point on them as the centre of rotation, the largest movement occurred between the different sets of 4 images.

A method was developed to exploit this to simplify the registration problem. The first step is to find this common rotation point on each of the 4 sets of images. This is done by finding the crossing point of the anterior segment from the 4 images. This point is defined as a minimum of the height difference between the different boundaries. This is done in turn for each set of 4 images and 4 different centres are found.

The largest motion occurs between different sets of 4 scans due to the increased time difference between the scans at this point in the imaging. If the relative position of these scans can be determined then this can be used to align the images. This is attempted using the ellipsoid approximation of the shape for the cornea. Once the centre of rotation for each image has been found this can then be used to generate four different sets of points to cover the cornea. For each set an ellipsoid is fit to the data using a least squares fitting method. These four ellipsoids should all correspond to the same surface. The centre of the different ellipsoids is found and all the points are adjusted to align the centre of the fitted ellipsoids.

When this step was tested with synthetic data that were generated by taking slices from an ellipsoid it worked well to align the data. When using real data the ellipsoids generated were not all the same so although this step led to an improvement in the alignment, it was not able to perfectly align real data. It was, however, used as an input for the GA program to reduce the range of possible values for the coefficients used.

4.4.3.3 Non GA methods

The non GA techniques evaluated are:

1. No registration
2. ICP
3. Height and centre from corner points, straight line version
4. Height and centre from corner points, curved version
5. Ellipse from corner points
6. Height and centre from pupil edges, straight line version
7. Height and centre from pupil edges, curved version
8. Ellipse from pupil edges

Each of these will be described in turn, and then the results of all the techniques will be compared. The goal of all the techniques is to combine 16 segmented 2D images to produce two 3D surfaces representing the anterior and posterior surfaces of the cornea.

4.4.3.4 No Registration

The simplest method of combining the images is to make the naïve assumption that the images are all perfectly aligned. The 16 images are assumed to all be from a series with different angles. The centre of each image, defined as the apex of cornea, is assumed to be a common rotation axis for all the 16 different images.

4.4.3.5 Iterative Closest Point

Iterative closest point (ICP) registration technique is a widely used image registration technique. It is designed to register two sets of points covering the same area. It works by calculating the minimum distance from each point in the reference set with the closest point in the target data set. An affine transformation is then found that moves the reference set of points to be closer to the target set. This process has repeated iterations in order to minimise the distance between the two sets of points.

The task faced here of registering a set of 16 images is not the same as the task ICP is designed for, since there are 16 images of different slices of the cornea with only a single line on them that corresponds to a second image. It is possible, however, to exploit the roughly symmetrical nature of the cornea to use ICP. Since each image is of a similar object, it is possible to align all 16 images to each other using ICP. The first image of the 16 can be considered as the template image and all 15 other images transformed to align with this image. The apex of the first image can then be calculated and assumed to be the centre of rotation about which all the images are varied. This allows a 3D surface to be constructed from the images.

4.4.3.6 Height and Centre from Corner Points, Straight Line Version

Since the segmentation of the images has been carried out it is relatively simple to find the location of various points of interest in the images. These points can be used to align the set of images. There are two obvious choices of points available to use. These are the corners of the anterior segment where the iris joins the cornea and the edges of the pupil. In addition to these points the apex of the cornea is also a useful point to define. These points are labelled in **Error! Reference source not found.** below.




Figure 43: Anterior segment OCT with key points for registration marked on it. The red points are the “corner points” where the iris and cornea join. The green points are the edges of the pupil. The pink point is the apex of the cornea.

The first attempt uses the corners of the cornea as reference points. The identification of these points has some slight difficulties due to the fact the cornea and the iris were not connected on all images. To get round this problem, a straight line approximation of the iris is used to find the corner points. A line ten pixels higher than the brightest horizontal line in the image is used to represent the top edge of the iris. The segmentation already includes a step to locate the iris using this method. Detection of the iris was carried out on a rotated image so it is able to detect the iris even when the image is not perfectly horizontally aligned. Where this line intersects with the lower boundary of the cornea is taken to be the corner points of the anterior segment. This was a reasonable approximation of the corner point over the set of images from 17 eyes.

Once the two corners were found the apex of the cornea is also located. This was taken as the point where the z value of the anterior surface is lowest (since $z = 0$ was defined at the top edge of the image). Using these three points a measurement of the anterior chamber height can be carried out. This is done by finding the mean z value of the two corner points and then finding the difference between that and the z value of the apex. This is done for the entire set of 16 images. There is a variation in the different heights found from the different images. The maximum anterior chamber depth is assumed to correspond to the image that went through the centre. The smaller heights from the other images are assumed to be due to the image not passing through the centre. Since all 16 images are of the same anterior segment the height for all of them should be the same.

The centre of each image is determined as the point midway between the two corners of the anterior segment. This point is assumed to be the point closest to the centre of the cornea with a distance from the centre d determined by

$$d_i = (c_0 - x_0) - \frac{h_i(c_0 - x_0)}{h_0} \quad (128)$$

where c_0 is location of centre in image with maximum anterior chamber depth, x_0 is location of corner of anterior chamber in same image, h_0 is anterior chamber depth for same image and h_i is anterior chamber depth of image i . This is a simple model and represents the cornea as a pyramid rather than the curved surface it has in reality.

The rotation in the plane of the image is determined from the angle the line between the two corners makes with the horizontal direction in the image. Displacement in the z direction is determined by aligning the mean z values of the corners of the anterior chamber.

4.4.3.7 Height and Centre from Corner Points, Curved Version

This method is very similar to the previously described method. The difference is the method of determining the displacement from perfect alignment using anterior chamber height. Since the previous assumption of a linear relation between the anterior chamber height and distance from centre is obviously not true, a different method was used. The image with the maximum anterior chamber height was used as a template. For each image the point on the template image that has the same anterior chamber height was found. The distance from this point to the centre of the image was taken as the distance from the image to the centre. The z displacement and angle was determined in the same way as the previous method. **Error!** **Reference source not found.** shows an example central image with the corresponding location of the centre of the other images marked on.




Figure 44: Image with upper surface marked on in red. This image had the largest anterior chamber height so is assumed to be crossing the centre of the image. The green asterisks demonstrate places where the height of each of the other images corresponds to the height in this image. The displacement of each image from centre is taken as distance from left most asterisk to each asterisk.

4.4.3.8 *Ellipse from Corner Points*

This method uses the corner points of the anterior segment, which are found the same way as described in the two previous methods. The difference is in how these points are used to determine the relative location of the images. For this method an assumption that the corner points of the anterior segment should form an ellipse is made. Since this corresponds to the iris being elliptical in shape it is true for all healthy patients and most diseased cases.

An ellipse is fitted to the half of the points with the largest distance between them; this is based on the assumption that these points with a greater separation have a smaller shift from the centre than those with a smaller separation. The distance required to shift each pair of points in order to align them with this ellipse is then calculated. This shift is used as the shift for all the points on the corresponding image allowing all the images to be aligned. A diagram of the ellipse and points before and after they have been fitted can be seen in **Error! Reference source not found.** below.




Figure 45: Diagram showing location of original corner points from all the images in blue, the ellipse that was fitted to the larger half of the blue points is shown in red. The points are then shifted to improve their fit to the ellipse, these points are shown in green.

4.4.3.9 Height and Centre from Pupil Edges, Straight Line Version

As an alternative to using the corner points of the anterior segment, the edges of the pupil can also be used. These are also found during the segmentation process.

The edges of the pupil are determined by first segmenting the pupil. The graph theory technique used means that only lines that go from one side of the image to the other are detected. The top surface and bottom surfaces of the iris are therefore segmented separately. The edges of the pupil are defined as the first and last points where the two lines are closer than 5 pixels. This is because where the pupil doesn't exist on the image segmentation of the top and bottom edges of the pupil will coincide with each other.

The location of the top of the anterior chamber was calculated using the same method as when using the corner points methods. Once the points have been

detected they can be used to register the images. This was done in the same way as when using the corner points from the cornea.

4.4.3.10 Height and Centre from Pupil Edges, Curved Version

This method is a combination of previously described techniques. It uses the pupil edges as described in the previous method. Then instead of using the straight line method to determine how this related to the displacement of the images, the curved distance method described in method 4 was used.

4.4.3.11 Ellipse from Pupil Edges

This method is also a combination of previously described techniques. This time the edge points of the pupil are used and the assumption is made that the pupil should be elliptical in shape. The same technique used in method 5 was used to fit an ellipse to the points and shift them accordingly.

4.4.4 Restricting the Range of Solutions in GA

It is possible to restrict the range of values that solutions that the GA can be allowed to take. The methods discussed in the previous section can be used to do this. A number of different methods were attempted. These all used the Zernike model energy function since this performed best.

4.4.4.1 Genetic Algorithm with Method 3 as Limit

In this method the step described in method 3 (corner points with straight line) is used as a first estimate. The available solutions are then limited to being near to the solution that is found from that. The fitness function used is based on using a model of what a cornea should look like. Each slice was free to move in 3 different directions and to rotate in the plane of the image.

4.4.4.2 Genetic Algorithm with Method 4 as Limit

Since the area of possible solutions is limited by the initial step which method is used to generate the initial estimate will affect the results. Therefore the GA method was also tried with a different method to find an initial estimate. Instead of using method 3 method 4 (corner points with curved line) was used to limit the range of results. Everything else was done in the same way.

4.4.4.3 Genetic Algorithm with Method 3 as Limit, X and Y Constraints Only

In the previous two methods attempted there are four degrees of freedom for each slice to move. Each of these different directions to move is constrained by the initial step. The determination of the angle and the vertical height are thought to be less critical than the x and y directions (i.e. movement in plane of image). The

estimation of these coefficients by the initial step is also less accurate. This method is therefore repeated but with the z and theta values having a much larger range of possible values than previously so they are not influenced by the initialisation.

4.4.4.4 Genetic Algorithm with No Limit

A test was carried out without using limits; previously solutions of the GA were limited to a region around the results of the initial estimate. This constraint is relaxed to enable the GA to find solutions over a much larger area. The same GA energy function is used in the minimisation.

4.4.5 Comparison of Different Methods

In order to test the different methods for registration a data set of 17 eyes was used. Each eye in this set has been imaged twice using the Visante OCT device. 16 radial scans were acquired for each scan. These images are then segmented using a graph theory based method. Using the segmented results the initialisation step based upon identifying the corners of the anterior chamber is used. This gives approximate values for shift needed for the image registration. Once this step is carried out the different energy functions for use with GA are tested.

The six different GA methods tested are circles in r, circles in r and z, ellipse fitting, ellipse fitting and circles, the model based method and the model based method with four rigid subsets of data. This is carried out for each of the two sets of images independently. This gives two surfaces that correspond to the corneal surface for 17 eyes. The distances between the two different surfaces of each of the eyes are calculated. These are used to give an unsigned mean distance between the two different surfaces. Since the same eye has been imaged twice, this distance should be zero. There will, in reality, be differences due to three different factors, the constancy of the OCT images (these could be caused by movement artefacts, limits to resolution etc.), the repeatability of the segmentation technique (some images may be slightly incorrectly segmented) and finally the errors may come from the registration program. The results of this comparison are shown in Chapter 5.

4.5 Conclusion

This chapter has described a number of different 3D segmentation and registration techniques. The effectiveness of these techniques has been evaluated and is described in Chapter 5.

5 Results

5.1 Introduction

Many different techniques for the segmentation and alignment of images have been developed and evaluated. This chapter will present the results of each technique in the same order they are introduced in previous two chapters. The first results that are discussed are from level set segmentation and a number of variations on this technique. Following this graph cut techniques will be discussed. The next section described graph theory based techniques. A comparison of the best of each category of 2D segmentation technique is then carried out. Discussion then moves onto 3D segmentation. Limited results from level set techniques are discussed. Following this the results of the graph cut techniques are discussed. Then the results of the different alignment methods used in conjunction with graph theory based segmentation are presented. Finally a comparison of all the different 3D segmentation techniques is carried out.

For the 2D techniques all results given are from comparison over a set of 39 images using 3 different comparison techniques. For the 3D techniques a set of images from 17 eyes was used for the comparison. The acquisition of the images is described in section 0 and the comparison techniques are described in section 1.6.

5.2 Level Set based Segmentation

This was the first method that was developed. There are a number of variations of this method that will be discussed in turn.

5.2.1 CVWS

The results of the segmentation technique described in Section 3.5 were compared with manual segmentation. A variety of different values were used for λ_3 which controls the strength of the shape term. The results are shown in Table 1 below. It can be seen from looking at the table that the best value for the shape term was 0.8 when looking over all the images. The optimum strength of the shape term varied from image to image. It can also be seen that for some images the value of the DSC is very low and the segmentation must be considered to have failed completely.

Table 1: Table showing results DSC comparison between manual segmentation and CVWS method.

λ_3	0.4	0.5	0.6	0.8	1
com43	0.69687	0.914702	0.919021	0.915554	0.917219

lss56	0.681387	0.693383	0.730572	0.96348	0.962151
wqm02	0.944704	0.941476	0.945606	0.942916	0.937919
wqm05	0.974852	0.975637	0.974791	0.968246	0.624081
wqm06	0.970688	0.964875	0.960259	0.938379	0.914459
wqm07	0.883905	0.928563	0.925306	0.91488	0.276295
wqm08	0.807417	0.914068	0.910326	0.904756	0.904766
wqm09	0.963421	0.960584	0.956396	0.933631	0.942253
wqm10	0.965128	0.96603	0.966753	0.962295	0.962295
wqm11	0.936854	0.931306	0.926313	0.91706	0.906015
wqm14	0.838199	0.967414	0.965677	0.955557	0.662707
wqm17	0.778664	0.814526	0.907893	0.877215	0.44921
wqm19	0.419427	0.405179	0.404299	0.395017	0.928888
wqm20	0.744319	0.927846	0.920874	0.907612	0.754848
wqm21	0.891554	0.907253	0.922712	0.912632	0.902397
wqm22	0.932438	0.929505	0.926084	0.92413	0.916456
wqm24	0.961191	0.958119	0.95454	0.919094	0.925107
wqm26	0.909235	0.90771	0.920619	0.95474	0.544088
wqm27	0.924672	0.93774	0.93423	0.932112	0.91656
wqm28	0.933809	0.926039	0.920143	0.90988	0.87332
wqm29	0.964532	0.962931	0.509411	0.156763	0.405487
wqm30	0.87589	0.926011	0.895854	0.888191	0.864667
wqm31	0.93398	0.934784	0.931193	0.934323	0.77395
wqm33	0.741178	0.808676	0.826225	0.939751	0.917182
wqm34	0.915255	0.912581	0.912348	0.910077	0.905411
wqm35	0.950037	0.949435	0.948234	0.942231	0.930831
wqm44	0.822617	0.830268	0.844166	0.944916	0.930557
wqm45	0.962168	0.957129	0.952508	0.938513	0.921794
wqm46	0.789332	0.831991	0.841203	0.928559	0.895899
wqm54	0.946889	0.949396	0.948218	0.943441	0.897521
wqm55	0.67367	0.794802	0.963106	0.958787	0.94103
wqm56	0.857215	0.873296	0.899428	0.943742	0.915563
wqm57	0.861785	0.896194	0.892862	0.896829	0.894033
wqm59	0.864574	0.955852	0.954302	0.928363	0.943474
wqm61	0.508391	0.470616	0.470634	0.426357	0.934428
wqm62	0.833174	0.858612	0.920089	0.920586	0.916128
wqm63	0.949797	0.950259	0.946041	0.940532	0.929915
wqm64	0.674105	0.812284	0.829337	0.92741	0.631691
wqm65	0.795838	0.892345	0.925028	0.945029	0.93845
Mean	0.848	0.883	0.882	0.884	0.839
Standard deviation	0.129	0.121	0.132	0.168	0.169

Boundary comparison techniques were also carried out for the images. In order to save space only the mean and standard deviation of these methods is shown in Table 2. Here it can be seen that the optimum value for the shape term varied depending on which technique is used to analysis the results. For the anterior

surface MSPE a shape constraint of either 0.5 or 0.6 performs best achieving a value of 1.66 pixels. For the posterior MSPE a shape constraint of 0.8 performs with a value of 2.89 compared to 4.52 when a shape weighting of 0.6 was used. With the Hausdorff distance measure the best results are achieved with a shape constraint of 1. **Table 2: Comparison between the CVWS automatic segmentation technique and manual segmentation using two comparison methods; the mean unsigned surface positioning error for both anterior and posterior surfaces and the Hausdorff distance measure.**

Value of shape term	0.4	0.5	0.6	0.8	1
MSPE (anterior boundary)					
Mean	1.75	1.66	1.66	1.77	4.05
Standard Deviation	1.25	1.24	1.23	1.34	5.69
MSPE (posterior boundary)					
Mean	6.38	5.00	4.52	2.89	5.46
Standard Deviation	3.89	3.63	3.21	1.33	5.87
Hausdorff distance					
Mean	28.5	21.0	19.5	16.0	15.2
Standard Deviation	35.5	34.1	33.9	31.0	22.8

The differences are most likely caused by the fact the Hausdorff distance is a measure find the greatest distance between the manual and automatic segmentation whilst the other techniques provide a mean measurement of difference. The differences are also likely to arise from the fact different parts of the image are segmented better with a different strength shape constraint. The posterior boundary is harder to segment so needs a stronger shape term to accurately identify the boundary whilst the anterior boundary is easier to detect and a stronger shape constraint could result in errors as the shape term is not a perfect model for the cornea.

Looking at these results it was decided that a value of $\lambda_3 = 0.8$ is the optimum value to be used for this method. Throughout the process the DSC measure of accuracy was used as the primary guide for strength of different terms since it is a technique that takes into account the entire image.




Figure 46: Example image showing both manual and automatic segmentation. The red line is automated segmentation using the CVWS technique. The Green line is manual segmentation by an expert. The image shown here is wqm35 which has a DSC value of 0.942




Figure 47: Example image showing two manual segmentations compared to automatic segmentation. The red line is automated segmentation using the CVWS technique. The Green line is manual segmentation by an expert. The blue line is manual segmentation by a second observer. The colours are mixed when the lines overlap

5.2.2 Comparison of Three Different Weighting Techniques

One way to improve the technique was to investigate alterations to the shape weighting technique as described in section 3.7.4. Three methods are compared, constant shape constraint across the image (CVWS), varying it with edge information (CVWSe) and varying it with intensity information (CVWSiw). Using iteration number to vary shape weighting produced obviously worse results in all cases so results were not included in this comparison.

In order to make the comparison the three different methods were carried out using a number of different values of λ_3 . This enabled the optimum value of λ_3 for each method to be found. The values of this were different for each method as is to be expected since the modifications to the code changed the value of the weighting term. The values of λ_1 and λ_2 were always kept the same. The optimum values of λ_3 used were 0.8 for the original constant method, 2 for the CVWSe method and 1.5 for the CVWSiw method. The results of comparison of the different methods are shown in Table 3 below.

Table 3: Table showing comparison between three different methods used to determine the weighting of the shape term

Method	CVWS	CVWSe	CVWSiw
DSC			
Mean	0.884	0.899	0.919
Standard Deviation	0.168	0.069	0.026
MSPE anterior			
Mean	1.77	2.26	1.64
Standard Deviation	1.34	1.72	0.52
MSPE posterior			
Mean	2.89	5.80	3.90
Standard Deviation	1.33	6.84	1.81
95% HD distance			
Mean	21.0	18.48	10.16
Standard Deviation	34.1	25.62	4.78

It can be seen that using the DSC comparison technique both methods varying the strength of the shape term across the image produced improved results. The CVWSe technique has DSC value of 0.899 and the CVWSiw technique has a value of 0.919 both of these are superior to the value of 0.844 achieved using the more basic CVWS technique.

Looking at the MSPE for anterior surface the CVWSiw again performs best with a value of 1.64 compared to 1.77 or 2.26 for the other two techniques. For the posterior surface the CVWS technique performs best with a MSPE value of 2.89 compared to 3.90 for the CVWSiw technique. The 95% HD distance technique results agree with the DSC values, the best result is produced with the CVWSiw technique achieving a value of 10.16 compared to 18.48 and 21.0 for CVWSe and CVWS techniques.

The method using the intensity weighting is the best technique of the three; the segmentation it produces is closer to the manual segmentation when using three of the four different techniques to measure it. It is worth noting that the original basic method has a much higher standard deviation than the other two techniques. This is because there are a number of images it completely fails on. If those images are removed the results for the method substantially improve and become comparable with the CVWSe method.




Figure 48: Illustration of agreement between the segmentations using the two new methods (CVWS and CVWSe) and the manual annotation. Red line is CVWS, green line is CVWSe and blue line manual annotation. Colours are mixed where lines overlap. Good agreement between the different methods can be seen especially on anterior surface.

5.2.3 Results of Texture Segmentation

The next technique that was evaluated used texture information instead of intensity to control the level set function. This is described in section 3.7.5. Table 4 shows the results of 4 different strengths of the shape constraint when comparison is made using DSC. From this the best results are achieved using a shape constraint of 0.4 with a DSC value of 0.909.

Table 4: Table showing results of DSC comparison between the results of manual segmentation and segmentation using the CVWSt method.

CVWS texture	0.2	0.4	0.6	0.8
com43	0.90615	0.894568	0.892336	0.856136
lss56	0.88031	0.923718	0.9566	0.958484
wqm02	0.909911	0.924832	0.911999	0.929377
wqm05	0.971598	0.963839	0.97116	0.963854
wqm06	0.965388	0.940672	0.950695	0.939927
wqm07	0.925217	0.897124	0.896296	0.900963
wqm08	0.878268	0.895005	0.8823	0.895931
wqm09	0.9565	0.918926	0.942245	0.913785
wqm10	0.92436	0.964288	0.928494	0.963769
wqm11	0.924739	0.912247	0.919234	0.91136
wqm14	0.795755	0.803231	0.737625	0.748837
wqm17	0.907575	0.908049	0.87851	0.875393
wqm19	0.904584	0.923821	0.942943	0.925399
wqm20	0.940795	0.956356	0.896606	0.862475
wqm21	0.847213	0.861009	0.886959	0.895195
wqm22	0.89905	0.897031	0.899558	0.901287
wqm24	0.948161	0.936476	0.94688	0.830741
wqm26	0.866941	0.87755	0.951287	0.867032
wqm27	0.901523	0.911434	0.907933	0.90784
wqm28	0.92335	0.907314	0.909335	0.90718

wqm29	0.861331	0.944825	0.661473	0.584269
wqm30	0.912185	0.906682	0.864903	0.807179
wqm31	0.899595	0.913818	0.91297	0.914791
wqm33	0.906436	0.920286	0.935194	0.916313
wqm34	0.860605	0.878484	0.903249	0.882086
wqm35	0.921236	0.92079	0.930137	0.935491
wqm44	0.827839	0.804783	0.825239	0.883594
wqm45	0.941907	0.941582	0.939615	0.941565
wqm46	0.918212	0.923085	0.921094	0.918701
wqm54	0.929779	0.927597	0.923516	0.92613
wqm55	0.893358	0.923298	0.948159	0.919534
wqm56	0.888992	0.87196	0.8892	0.88936
wqm57	0.912469	0.858536	0.851476	0.859515
wqm59	0.941446	0.943966	0.915379	0.941867
wqm61	0.87773	0.916385	0.923732	0.913928
wqm62	0.910772	0.889329	0.902412	0.888114
wqm63	0.938086	0.920379	0.917658	0.919629
wqm64	0.899685	0.921625	0.914995	0.900769
wqm65	0.908947	0.919875	0.9254	0.923884
Mean	0.906	0.909	0.903	0.893
Standard Deviation	0.036	0.035	0.057	0.066

It can be seen from this that using the CVWSt technique the results are optimised with a smaller value of the shape coefficient than the basic CVWS method. There is also a smaller standard deviation indicating fewer outlying results. Table 5 below shows the results using four different values of the shape constraint using the boundary comparison techniques. Again it can be seen that the different comparison techniques have different optimum results. The best MSPE value for the anterior surface was 2.66 which was found when a shape constraint of 0.6 was used. The best MSPE for the posterior surface is 16.92 which corresponded to a shape constraint of 0.2. The best 95% HD distance was 8.82 and this was achieved using a shape constraint of 0.8.

Table 5: Table showing comparison of segmentation using CVWSt and manual segmentation using three different line based methods. This first two are mean surface positioning error (MSPE) calculated separately for anterior and posterior boundaries, the third measure is the 95% Hausdorff distance calculated over the entire boundary.

Strength of shape constraint	0.2	0.4	0.6	0.8
MSPE anterior				

Mean	3.15	2.82	2.66	2.82
Standard Deviation	2.33	2.13	2.00	1.79
MSPE posterior				
Mean	16.92	17.45	18.13	18.77
Standard Deviation	7.32	7.56	7.58	7.81
95% HD distance				
Mean	24.17	15.21	10.10	8.82
Standard Deviation	33.34	26.27	12.92	6.06

The anterior boundary can be found more accurately than the posterior boundary. This is common to all methods as the anterior boundary is much better defined on the image. It was decided to use a weighting of 0.4 for the shape term since that gives the best value when using the DSC comparison method.

5.2.4 Dynamic Weighting and Texture

An investigation was also carried out to see how the inclusion of dynamic weighting would affect the performance of the texture model. Dynamic weighting based on the number of iterations and based on edge information from across the image was attempted.

As was found when using intensity information from the image segmentation attempts at altering the strength of the shape term by iteration number produced significantly worse results. A full analysis of these results was not carried out and all further work using those methods was abandoned. The results using edge information to vary the image intensity were more promising and are shown below.

Table 6 shows the results of this comparison between segmentation with the automatic method CVWSte and manual segmentation. The best DSC results of 0.89 are found when using a shape constraint of 0.4. This is repeated with the MSPE of the anterior surface with a value of 3.97. Looking at the posterior surface the best value was found using a shape constraint of 0.8. For the 95% HD distance the best value was found when a value of 0.8 was used for the shape constraint.

Table 6: Table showing comparison between segmentation using the CVWSte method and manual segmentation using different comparison techniques and different strengths of the shape constraint.

Strength of shape constraint	0.4	0.6	0.8	1

DSC				
Mean	0.890	0.889	0.889	0.883
Standard Deviation	0.041	0.051	0.059	0.070
MSPE anterior				
Mean	3.97	4.30	4.60	4.77
Standard Deviation	2.87	3.42	4.25	5.07
MSPE posterior				
Mean	17.60	16.91	16.81	16.84
Standard Deviation	7.48	7.29	7.05	7.05
95% HD distance				
Mean	26.5	25.0	22.4	26.1
Standard Deviation	38.3	41.8	24.5	26.4

Looking at these results a reasonable level of agreement with manual segmentation can again be seen. The same optimum weighting of the shape term is found as the method using the constant weighting. A comparison of the results of the two texture based methods shows they produce similar results.

5.2.5 Comparison of Different Level Set Methods

So far in this section five different versions of the novel level set with shape method have been presented. Two previously presented segmentation techniques were also implemented for the purpose of comparison. The five different versions are Chan and Vese with shape prior (CVWS), Chan and Vese with shape prior and dynamic weighting based on edge (CVWSe), Chan and Vese with shape prior and dynamic weighting based on intensity (CVWSiw), texture based Chan and Vese with shape prior (CVWSt) and texture based Chan and Vese with shape prior and dynamic weighting based on edge (CVWSte). The previously used methods are Chan and Vese's active contour without edges (CV) [70] this is equivalent to the techniques we have developed with the shape term $\lambda_3 = 0$ and a threshold based method developed by Shen et al. [129]. Results of manual segmentation by a second observer are also compared.

In order to carry out this analysis the results of segmentation for the 5 new models, two existing models and a second manual observer were compared to the segmentation by a manual observer. The results of these comparisons using DSC is shown in Table 7 below, the results using the line based methods are shown in Table 8.

Table 7: Mean results of DSC comparison between seven automatic segmentation techniques and manual segmentation. 39 images have been used for the comparison. Results are also shown of the difference between manual segmentation by two different experts.

Technique	Mean	Standard Deviation
CVWS	0.884	0.168
CVWSe	0.899	0.069
CVWSiw	0.919	0.026
CVWSt	0.909	0.035
CVWSte	0.89	0.041
CV	0.654	0.049
Threshold	0.767	0.1
Inter Observer	0.966	0.007

Table 8: Mean results of line comparison methods between six automatic segmentation techniques and manual segmentation over 39 images. This first two comparison techniques are mean surface positioning error (MSPE) calculated separately for anterior and posterior boundaries, the third comparison technique is the 95% Hausdorff distance calculated over the entire boundary.

	MSPE anterior		MSPE posterior		95% HD distance	
	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation
CVWS	1.87	1.27	16.1	5.7	34.3	66.3
CVWSe	2.26	1.72	16.6	6.0	49.1	78.6
CVWSiw	1.63	0.52	3.9	1.81	10.2	4.8
CVWSt	2.82	2.13	17.4	7.6	15.2	26.3
CVWSte	3.97	2.87	17.6	7.5	35.6	32.5
CV	3.21	1.59	12.46	2.36	20.07	19.34
Threshold	2.82	5.33	11.05	3.66	25.39	8.89
Inter Observer	1.07	0.21	1.16	0.34	36.4	7.3

Looking at the tables it can be seen that the newly developed methods have superior performance to previously reported methods. The best new method is the CVWSiw technique which achieved a DSC value of 0.919 compared to a value of 0.767 for the threshold method which is a previously published technique. This pattern is repeated in MSPE measurements with CVWSiw getting values of 1.63

and 3.90 for anterior and posterior surface respectively. This is better than the 2.82 and 11.05 values achieved using threshold. The results do not perform quite as well as manual observation which achieved a DSC value of 0.966 and MSPE for anterior and posterior surfaces of 1.07 and 1.16. This indicates the difference between two manual observers is less than the difference between manual observers and automated segmentation methods.

ANOVA analysis was carried out on the DSC results to determine if the differences between the methods were significant. There was a significant difference between inter observer difference and all the other methods ($p < 0.05$). There was also a significant difference between CV, the threshold technique, and all the other methods ($p < 0.05$). The different methods developed here showed less significant differences between them. A significant difference was found between CVWS and CVWSiw ($p = 0.04$). The differences between the other methods were not found to be significant at the 0.05 level. This can be seen by looking at the very close values of the mean DSC value for the different methods.

Figure 49 below shows a binary representation of the segmentation produced by the various different techniques. It can be seen by comparing (e) to images (a-d) the effect of the shape term. The regions in the middle with a lower signal to noise ratio can't be detected without the addition of the shape term which does detect them. The thresholding method shown in (f) avoids this problem by fitting polynomials to the surfaces of the cornea. While this achieves reasonable results the drawbacks can be seen in the bulge on the right of the image which is an artefact produced by the polynomial. Looking at images (g) and (h) good agreement can be seen between the best automatic technique and manual segmentation. The effect of the intensity weighting term can be seen in image (h), the posterior surface is very smooth indicating that the shape term has been allowed to dominate here due to the weak signal from the image.




Figure 49: Montage of results of segmentation by different methods. (a) is CVWS, (b) CVWSe, (c) CVWSt, (d) CVWSte, (e) CV, (f) threshold, (g) manual 1, (h) CVWSiw.

Figure 50 below shows segmentation by three different methods overlaid on the same image. It shows reasonable agreement between the methods as expected. The differences between the three methods are more pronounced in regions where segmentation is harder. It also gives an indication of why using dynamic weighting does not results in significantly improved results as might be expected, improving the segmentation in outer regions of higher signal results in a shape term that is less effective at modelling the lower signal regions pulling the contour too far away from the true boundary.




Figure 50: Example image with results of segmentation overlaying on it. Red line is CVWS method. Green line CVWSe. Blue line CVWSt. Where lines overlap the colours are mixed.

5.2.6 Discussion of Level Set Work

A fully automatic technique has been developed that can detect both the anterior and posterior surfaces of the anterior segment in AS OCT images. The algorithm used a shape prior to allow difficult-to-segment regions to be segmented. The technique has been demonstrated to be able to successfully segment images including regions with a low signal to noise ratio. The newly developed method performed significantly better than previously published methods and the results show a high level of agreement with expert manual segmentation.

A comparison between the different variations of the model has revealed there are no significant differences between the results of the different models. The CVWSiw model performed best with a higher mean value DSC than the others. This model is therefore taken as being the best of the options that were developed. The lack of significant differences is likely due to the relatively small sample size of the data and the similarity of the different methods.

One of the important factors that can affect the segmentation performance is the image quality such as the signal noise ratio (SNR). In general, the higher SNR an image has, the easier segmentation will be. For this particular problem the image contains speckle noise inherent in the OCT system and poor SNR in some the cornea structures. This means simple thresholding and region-based models will not work; this was demonstrated in the comparative study. Our method uses a shape term to overcome this problem.

Overall, the level set work has shown that using a shape prior term can significantly improve segmentation results of fully automatic segmentation of the cornea in AS-OCT images. This section of work has been published in the Journal of Biomedical Optics [1].

5.3 Graph cut

The next method attempted was a graph cut based implementation of the Chan and Vase with shape model; this is described in section 3.8. Two different methods of representing the shape term were used and have been compared in turn.

5.3.1 Binary Shape model

The first technique developed was based on using a binary shape term. A comparison was carried out between the results of segmentation using graph cut and manual segmentation using the same method as previously described. The

results of this segmentation method carried out over the entire data set are shown in Table 9 below. Several different values of the shape constraint are shown in the table. The best DSC value of 0.894 is found using a shape term values at 0.04 and 0.05.

Table 9: DSC comparison between segmentation results of graph cuts program with the manual segmentation over all the images.

Graph Cut	0.03	0.04	0.05	0.06	0.07
com43	0.869278	0.872199	0.879135	0.883076	0.895932
lss56	0.906831	0.900563	0.927237	0.935088	0.938201
wqm02	0.880874	0.885133	0.895797	0.905248	0.915779
wqm05	0.954421	0.954184	0.953437	0.951338	0.819703
wqm06	0.957545	0.953998	0.950696	0.942788	0.935481
wqm07	0.899124	0.901237	0.902056	0.901335	0.900501
wqm08	0.874976	0.882752	0.876393	0.87382	0.870847
wqm09	0.955497	0.946818	0.935639	0.923529	0.897523
wqm10	0.925368	0.929992	0.931801	0.931213	0.925014
wqm11	0.919018	0.918529	0.916267	0.910018	0.897051
wqm14	0.591719	0.602654	0.591449	0.586053	0.584795
wqm17	0.871978	0.868935	0.866685	0.860662	0.858766
wqm20	0.915167	0.906104	0.888573	0.841816	0.807155
wqm21	0.848768	0.849427	0.853126	0.854586	0.854586
wqm22	0.898239	0.882051	0.888646	0.886537	0.891197
wqm24	0.948845	0.948212	0.947136	0.94502	0.940229
wqm26	0.955488	0.951725	0.948585	0.941928	0.937546
wqm27	0.51551	0.900599	0.898562	0.901436	0.899093
wqm28	0.905915	0.91037	0.908877	0.90874	0.904571
wqm29	0.810612	0.810612	0.810612	0.810612	0.810612
wqm30	0.871544	0.873364	0.869619	0.867921	0.862609
wqm31	0.900328	0.899556	0.901487	0.900596	0.90293
wqm33	0.519375	0.90119	0.903361	0.910795	0.910782
wqm34	0.85318	0.840745	0.840975	0.846644	0.846642
wqm35	0.915663	0.917667	0.919444	0.918077	0.919178
wqm44	0.516463	0.910522	0.919469	0.919335	0.927338
wqm45	0.937698	0.930417	0.93188	0.931612	0.93047
wqm46	0.920393	0.9126	0.91164	0.911919	0.912536
wqm54	0.569336	0.913263	0.913636	0.913231	0.913607
wqm55	0.929903	0.935039	0.936934	0.934914	0.936619
wqm56	0.564787	0.928972	0.929567	0.929803	0.929584
wqm57	0.579056	0.846376	0.842659	0.844899	0.834986
wqm59	0.951669	0.951784	0.949913	0.94874	0.946123
wqm61	0.51762	0.891252	0.872027	0.874734	0.877477
wqm62	0.577575	0.8507	0.846281	0.871465	0.872852
wqm63	0.607962	0.914107	0.911867	0.91521	0.912965

wqm64	0.531254	0.897457	0.894513	0.899418	0.900215
wqm65	0.903902	0.89881	0.908422	0.908708	0.909743
Mean	0.805	0.894	0.894	0.893	0.888
Standard Deviation	0.166	0.060	0.061	0.061	0.063

Table 10 shows the result of comparison using the line based comparison techniques. Again the optimum strength of the shape term varies with the method used for evaluation. Looking at MSPE for anterior and posterior the best results (4.08 for anterior and 4.55 for posterior) are found when a shape constraint of 0.07 is used. The best result for 95% Hasudorff distance 10.21 is found using a shape term valued at 0.06.

Table 10: Line comparison between segmentation results from graph cut with manual segmentation results.

MSPE- anterior	0.03	0.04	0.05	0.06	0.07
Mean	4.47	4.453	4.30	4.15	4.08
Standard Deviation	4.89	4.21	4.13	4.11	4.09
MSPE - Posterior					
Mean	4.89	4.98	4.80	4.64	4.55
Standard Deviation	5.01	4.32	4.25	4.22	4.21
95% HD					
Mean	32.92	11.14	10.39	10.21	10.38
Standard Deviation	39.00	7.19	6.06	6.10	6.13

The graph cut method performs slightly worse than the CVWS method it is based on. Looking at the line comparison it can be seen that unlike all the previous methods the results are just as good for the anterior and posterior surfaces. Both of these differences are probably due to the fact that a binary shape term has been used. This gives a more rigid shape constraint compared to the level set method. Our method of selecting the shape term through fitting ellipses produces a shape that is close to the cornea but is not an exact model.

Table 11 compares the time taken for the graph cut and level set segmentation techniques. The time taken to run the graph cut program is 4.40s which is significantly shorter than the level set based technique which took 151s. There is a significant reduction in segmentation time by a factor of over 30. The reduction in time taken is less pronounced than previous studies have found [130]. This is

probably due to the repeated calculation of the shape term that is carried out in our method.

Table 11: Table comparing mean time taken for segmentation using graph cut method (GC) and level set method (LS). The graph cut method shows a significant reduction in time needed for segmentation.

Technique	Graph Cut	Level Set
Mean time (s)	4.399	150.694
Standard Deviation	0.95	58.518

Overall the results of this technique were promising. A significant reduction in speed was achieved. The quality of the results is slightly worse using graph cut when compared to level set based methods. It was therefore decided to continue investigating this problem and alternative ways of implementing a shape term in graph cut segmentation were investigated. Figure 51 shows an example image segmented using this technique.




Figure 51: Example image of graph cut segmentation overlaid on original image. This is same image that has been show previously

5.3.2 Distance Function Shape Model

In order to further develop the model an approach was developed based on using a distance function rather than the binary shape term used in the previous model. The results of comparisons of this technique with manual observers are shown in Table 12. This technique achieves a DSC value of 0.943 when a shape constraint value of 0.1 was used. The best results for anterior MSPE, a value of 1.18, were found using a slightly smaller shape constraint of 0.05. The best results for MSPE posterior, 2.84, and 95% Hausdorff distance, 6.73, were found using a shape weighting term of 0.1.

Table 12: Results of comparison of graph cut segmentation technique with manual segmentation at different strengths of the shape term.

Value of λ_{shape}	0.01	0.05	0.1	0.2	0.3	0.4
DSC						
Mean	0.522	0.873	0.943	0.941	0.936	0.932
Standard Deviation	0.114	0.135	0.020	0.020	0.021	0.024
MSPE- anterior						
Mean	1.82	1.18	1.21	1.86	2.40	2.89
Standard Deviation	1.17	0.31	0.31	1.84	2.99	3.94
MSPE - Posterior						
Mean	4.55	4.95	2.84	2.86	3.13	3.43
Standard Deviation	3.23	4.17	1.278	1.39	1.70	2.18
95% HD						
Mean	103.82	19.20	6.73	6.92	7.63	8.30
Standard Deviation	28.41	26.81	2.65	3.51	5.02	6.07

The DSC values indicate that the segmentation performs badly if a weaker shape term than 0.05 is used. The results for the anterior surface are best at a value of 0.05. With an increasing shape term the location of the anterior surface gets worse. The relatively good value of 1.82 when a shape weighting term of 0.01 is used compared to the very poor result when using DSC of 0.522 is due to a weakness of the MSPE technique. The surface positioning error can only be defined on surfaces that exist. When the segmentation performed very badly this will leave gaps in the surface, these have been ignored giving a better result that is warranted by the segmentation. The 95% Hausdorff distance also shows substantially worse when a small weighting term is used, 103.8 for a shape term of 0.01.

Based on these results it was decided to fix the shape weighting at 0.1 for all future analysis work. This was because this performed better than using a weighting of 0.2 on all of the analysis methods. The differences between using a weighting of 0.1 and 0.2 is not statistically significant. It is possible that with a larger sample size or a different data set a different optimum value might be found. However the small difference between the two options means this difference is unlikely to be large.

Table 13: Comparison of time taken for two variations of graph cut method and level set method.

Time (s)	GC(v1)	GC(v2)	LS
Mean	4.40	3.53	150.7
Standard Deviation	0.95	0.63	58.5

Table 13 shows a comparison of the time taken to segment a single image by the two graph cut methods and the level set technique. The time taken for the new graph cut method GC(v2) is similar to the other graph cut method, 3.53s compared to 4.40s. This is to be expected as only a small change has been done to the method. The difference in time between the two versions of the graph cut program is not statistically significant.

5.3.3 Discussion of Graph Cut Segmentation

Two automatic segmentation techniques using graph cut techniques have been implemented. Both of these new techniques have produced a significant reduction in the time taken for segmentation in relation to the level set technique. The second technique has demonstrated an increased accuracy in relation to the level set technique achieving a DSC value of 0.943 compared to 0.919 that can be achieved using the best of our level set techniques. This shows that this graph cut technique is superior to the level set technique in relation to both speed and accuracy.

A paper presenting this technique has been accepted for publication by Eye and Vision Science journal.

5.4 Graph Theory Segmentation

Graph theory segmentation was investigated to see if this method could achieve further improvements in speed and accuracy of segmentation. Two variations of graph theory segmentation were attempted. The first was based on using Dijkstra's algorithm to find the shortest path; the second method instead used dynamic programming to find the shortest path. The second method has been implemented twice, first using Matlab and then with C++, the two versions of the code are compared here. The technique is described in section 3.9.

5.4.1 Dijkstra's Algorithm Method

The Dijkstra's algorithm based method was used to segment the same data set as the previous techniques and the segmentation results were compared with manual segmentation results. A number of different strengths for the shape coefficient were tested in order to find the optimal value. The results are shown in Table 14 below.

Table 14: Results of Graph Theory segmentation and DSC comparison with manual observers

	0.3	0.4	0.5	0.6	0.7	0.8
com43	0.906593	0.906593	0.907006	0.907073	0.951681	0.951398
lss56	0.962526	0.962655	0.962474	0.962474	0.962474	0.962517
wqm02	0.959999	0.959666	0.959683	0.959542	0.959469	0.959526
wqm05	0.967892	0.96791	0.96791	0.968509	0.968509	0.967615
wqm06	0.971572	0.971572	0.971572	0.971572	0.971572	0.971534
wqm07	0.953963	0.954362	0.954362	0.954362	0.954362	0.954262
wqm08	0.932141	0.932257	0.932257	0.932241	0.932357	0.932631
wqm09	0.917199	0.917749	0.922892	0.913356	0.911104	0.911104
wqm10	0.767784	0.959614	0.959909	0.960037	0.960037	0.959777
wqm11	0.962834	0.963363	0.963363	0.963363	0.963363	0.963363
wqm14	0.964182	0.964182	0.964537	0.964537	0.96429	0.964332
wqm17	0.954913	0.954785	0.954785	0.954785	0.954748	0.954821
wqm19	0.945502	0.945525	0.94563	0.94563	0.945543	0.945473
wqm20	0.952465	0.952465	0.952835	0.952835	0.953419	0.953419
wqm21	0.958077	0.958077	0.958077	0.958129	0.957496	0.957496
wqm22	0.947495	0.947339	0.947066	0.946861	0.946861	0.946707
wqm24	0.953639	0.953119	0.953119	0.952927	0.952979	0.952979
wqm26	0.967377	0.967253	0.967253	0.967253	0.967213	0.966835
wqm27	0.948631	0.948631	0.948346	0.948346	0.948346	0.948331
wqm28	0.960055	0.960055	0.960055	0.960724	0.960724	0.960724
wqm29	0.960453	0.960453	0.960453	0.960453	0.960453	0.960856
wqm30	0.949601	0.949619	0.949967	0.949601	0.951977	0.951977
wqm31	0.96005	0.959767	0.959767	0.959767	0.959767	0.959767
wqm33	0.958383	0.958284	0.958502	0.958502	0.958502	0.958502
wqm34	0.768678	0.773111	0.949141	0.949062	0.933117	0.931709
wqm35	0.959503	0.959503	0.959503	0.959503	0.959215	0.959215
wqm44	0.955242	0.955242	0.955242	0.955242	0.955242	0.955242
wqm45	0.955923	0.955923	0.9565	0.955967	0.955967	0.955967
wqm46	0.951844	0.951844	0.951844	0.951761	0.951678	0.951678
wqm54	0.932919	0.956246	0.956246	0.956269	0.956542	0.956487
wqm55	0.966037	0.966037	0.966037	0.96602	0.966054	0.96602
wqm56	0.956593	0.956593	0.956436	0.956548	0.95649	0.956294
wqm57	0.920934	0.920726	0.928497	0.927262	0.927262	0.927367
wqm59	0.961115	0.961338	0.961242	0.961242	0.961179	0.961148
wqm61	0.908624	0.960235	0.959756	0.958053	0.957763	0.957023
wqm62	0.915137	0.922463	0.954023	0.954036	0.954084	0.951085
wqm63	0.965509	0.965509	0.964412	0.964412	0.964026	0.964026
wqm64	0.953463	0.953463	0.953346	0.953187	0.947875	0.947663
wqm65	0.958907	0.958907	0.959781	0.814183	0.833285	0.833285
Mean	0.941	0.949	0.954	0.95	0.951	0.951
Standard deviation	0.044	0.032	0.013	0.026	0.023	0.023

The best value for the strength of the shape term varied between images but a value of 0.5 produced the best result for the entire data set, this gave a mean DSC value of 0.954. This has both the highest mean value 0.954, and the lowest standard deviation 0.013, indicating the results are all clustered around a similar region to the manual segmentation. An example image segmented by this technique is shown in Figure 52 below.




Figure 52: Example image showing OCT image with graph theory segmentation path marked on it.

Table 15 shows the results using the line based comparison techniques. The same values of the shape constraint were used. Here different optimum values were found. Looking at MSPE for the anterior surface the changing value of shape constraint had very little effect this is to be expected since the segmentation of this surface is not affected by the shape term. The best value achieved was 1.16. Looking at the posterior surface the smallest MSPE 1.77 was achieved when a shape value of 0.7 was used. This strength of the shape term also produced the best 95% Hausdorff distance of 4.18.

Table 15: Comparison between automatic segmentation using graph theory minimised by Dijkstra's algorithm and manual segmentation using line based comparison techniques.

Value of λ_{shape}	0.3	0.4	0.5	0.6	0.7	0.8
MSPE - anterior						
Mean	1.16	1.16	1.21	1.21	1.16	1.21
Standard Deviation	0.29	0.29	0.46	0.46	0.29	0.46
MSPE - posterior						
Mean	2.64	2.15	1.93	1.85	1.77	1.85
Standard Deviation	2.73	1.87	1.39	1.14	0.77	1.15
95% HD						
Mean	6.88	5.31	5.07	4.66	4.18	4.67
Standard Deviation	7.59	5.28	5.90	5.09	2.92	5.08

5.4.2 Dynamic Programming Method

The second method involving graph theory used dynamic programming rather than Dijkstra's algorithm. The results achieved using the dynamic programming method implemented using Matlab are shown in Table 16 below.

Here the best results are achieved using a shape constraint of 0.001 for all comparison techniques. The highest DSC value found was 0.960, no difference was found between using a shape term of 0.0005 and 0.001 using this technique. The anterior surface location was the same regardless of the value of the shape constraint as the location of this surface is not affected by the shape term, a MSPE of 1.17 was found for the anterior surface. For the posterior surface the best MSPE was 1.49. The lowest 95% HD distance was 3.37.

A comparison was also made to the method with no shape term being used. Here it was found that the results were significantly worse when looking at the posterior surface. The anterior surface is not affected by the shape term so is unchanged. The DSC value found with no shape was 0.668 much lower than the 0.960 found with a shape term. The MSPE of the posterior surface was 17.3 indicating this was the surface that the segmentation technique was unable to find. This comparison was included since it is similar to methods previously published by others [97].

Table 16: Comparison between automatic segmentation using graph theory minimised by dynamic programming and manual segmentation.

Value of λ_{shape}	no shape	0.0001	0.0005	0.001	0.005
DSC					
Mean	0.668	0.958	0.960	0.960	0.954
Standard Deviation	0.125	0.010	0.009	0.009	0.012
MSPE – Anterior					
Mean	1.17	1.17	1.17	1.17	1.17
Standard Deviation	0.37	0.37	0.37	0.37	0.37
MSPE – Posterior					
Mean	17.3	1.80	1.60	1.49	1.75
Standard Deviation	4.67	0.59	0.45	0.47	0.59
95% HD					
Mean	28.5	4.40	3.80	3.37	4.21
Standard Deviation	2.9	1.64	1.20	0.89	1.50

These results indicate that the dynamic programming results are marginally better of the two implementations, achieving a DSC value of 0.960 compared to 0.954. The difference between the two implementations of the technique is very small and is not statistically significant. The much lower strength of the shape term used with dynamic programming is not significant just a reflection of the slightly different way the shape term was generated in the different versions of the method.

5.4.2.1 Dynamic Programming using C++

This same technique was implemented using C++ instead of Matlab to try to reduce the running time. Table 17 shows the results of this method. Here the results are slightly different to the previous implementation of the same technique. The best results were achieved using a shape value of 1. The best DSC value found was 0.962 which is marginally better than the best value achieved using the Matlab implementation of the same method, which found a value of 0.960. As before the MSPE for the anterior surface is unchanged by shape, a value of 1.10 was found. For the posterior surface the best value was 1.75 pixels. The best value for the 95% Hausdorff distance was 4.10. The results of this method were similar to the implementation of the method using Matlab. The difference between the two implementations is not statistically significant.

Table 17: Results of comparison between manual segmentation and automatic segmentation using graph theory and dynamic programming implemented in C++

Value of λ_{shape}	0.2	0.4	0.7	1	2
DSC					
Mean	0.950	0.957	0.961	0.962	0.960
Standard Deviation	0.014	0.011	0.010	0.009	0.011
MSPE – Anterior					
Mean	1.10	1.10	1.10	1.10	1.10
Standard Deviation	0.33	0.33	0.33	0.33	0.33
MSPE – Posterior					
Mean	2.56	2.17	1.81	1.75	1.80
Standard Deviation	1.00	0.90	0.69	0.60	0.76
95% HD					
Mean	7.04	5.49	4.23	4.10	4.17
Standard Deviation	3.22	2.70	1.85	1.61	1.74

5.4.3 Graph Theory Time Comparison

Three versions of graph theory segmentation have been implemented. Graph theory with Dijkstra's algorithm, graph theory with dynamic programming in Matlab and graph theory with dynamic programming in C++. The time taken for each of the different methods is shown in Table 18.

The fastest technique is the graph cut with dynamic programming implemented in C++ this has a mean time of 0.288s to segment a single image. This is much faster than the 16.1 s that the Matlab implementation of the same method took. The method using Dijkstra's method was the slowest technique taking 125.1s to segment a single image. Note that the time for the dynamic programming includes

time taken to identify boundaries of the iris. The time taken could be further reduced by skipping those steps.

Table 18: Time taken for segmentation of the three different implementations of the graph theory technique

Technique	Mean Time (s)	Standard Deviation
Graph theory Dijkstra's algorithm	125.1	13.1
Graph theory dynamic programming (Maltab)	16.1	0.8
Graph theory dynamic programming (C++)	0.288	0.013

Overall looking at the different graph theory methods the technique using dynamic programming performs best. It is faster than the Dijkstra's technique and has a marginally improved accuracy. Implementing this method using C++ managed to reduce the time taken for segmentation by a factor of 55 without any loss in accuracy of the technique. It is likely other techniques could also be made faster through a similar implementation in C++. This was not done since the Graph Theory technique was the one that produced the best results.

5.5 Comparison of Different Segmentation 2D Techniques

The segmentation techniques described here all fall into 4 main categories. In order to evaluate the methods developed the best method from each of these categories were compared. These results were compared with manual segmentation results. The four different methods to be compared are; the threshold based technique, Chan Vese with shape technique (the best of the level set methods), our graph cut implementation of Chan Vese with shape, and the dynamic programming implementation of graph theory segmentation. These methods were all applied to the same set of 39 images. These images had been segmented manually by two different observers and the first observer had segmented the images twice independently. This allows us to make a comparison with inter and intra observer variation.

Three different comparison techniques were used to look at the images. These are the Dice Similarity Coefficient (DSC), unsigned mean surface positioning error (MSPE) and 95% Hausdorff distance. The results for the different methods will be presented in turn.

Table 19: DSC comparison between four different automatic segmentation techniques. All techniques have been compared against the same set of manual annotations over a set of 39 images.

Segmentation Method	Mean	Standard Deviation
Threshold	0.852	0.082
CVWSiw	0.919	0.026
Graph Cut	0.943	0.020
Graph Theory	0.962	0.009
Inter Observer	0.966	0.007
Intra Observer	0.969	0.010

Table 19 above shows the results of the DSC comparison with manual segmentation for the four different automated methods and two other manual segmentation results. In order to test the significance of the results an ANOVA test was carried out using SPSS software. This showed that there was a significant difference between the different methods. Doing a Turkey post hoc analysis found there are four overlapping homogenous subsets of the techniques. The Threshold method is significantly different to all the other methods ($p < 0.001$ for all other techniques), achieved a DSC value of 0.852. Using CVWSiw a DSC value of 0.919 was found, this is significantly different to all the other methods ($p < 0.05$ for all other techniques). The graph cut technique achieved a value of 0.943, this was not significantly different to either the graph theory technique ($p = 0.270$) or the inter observer difference ($p = 0.056$). There was a significant difference with the intra observer difference ($p < 0.05$). The lack of significance of this difference is probably caused by the sample size rather than the technique. When an independent t test looking at just the inter observer difference and graph cut technique was carried out it did show a significant difference between these techniques. The Graph Theory technique, Inter Observer and Intra Observer differences achieved the best results getting DSC values of 0.962, 0.966 and 0.969 respectively. There was no significant difference between these results ($p = 0.950$ as a grouping).

In summary the DSC results indicate that all three techniques that have been developed are better than a previously used technique (threshold). The graph theory based technique was the best performing method developed. The graph theory segmentation technique is able to perform equally as well as manual observers. There is no significant difference between inter and intra observer segmentations indicating that good quality objective segmentation of these images is possible and results aren't dependant on the person carrying out the imaging.

Table 20: MSPE comparison between four automatic segmentation techniques and three manual segmentations. The anterior and posterior surfaces have been compared separately.

Segmentation Method	Anterior Surface		Posterior Surface	
	Mean	Standard Deviation	Mean	Standard Deviation
Threshold	1.98	4.98	11.05	3.66
CVWSiw	1.63	0.52	3.9	1.81
Graph Cut	1.21	0.31	2.84	1.28
Graph Theory	1.10	0.33	1.75	0.60
Inter Observer	1.31	0.53	1.78	0.67
Intra Observer	1.08	0.33	1.32	0.50

Table 20 above shows the results of mean surface positioning error test comparing the four different automated methods with manual segmentation results. An ANOVA test was carried out using MSPE measurements from the anterior and posterior surfaces separately. The results for this test on the anterior surface were less conclusive than the DSC analysis. All the results have an anterior MSPE of between 1 and 2 pixels. Of the automated techniques the threshold technique has the largest difference of 1.98. The smallest MSPE was found using the graph theory technique at 1.10 pixels, this is better than the inter observer MSPE of 1.31 but worse than the intra observer MSPE of 1.08. There were no statistically significant differences between the different techniques looking at the anterior surface.

Looking at the posterior boundary similar results to the DSC comparison were found. Analysis of the results indicates four overlapping subsets of results. The threshold technique performs significantly worse than all the other techniques ($p < 0.05$) with a MSPE of 11.05. The graph cut and level set techniques form a second group significantly different to the others ($p = 0.252$ between them), achieving a posterior MSPE of 3.9 for the level set technique and 2.84 for the graph cut technique. The graph cut did not show a significant difference when compared to either the inter observer difference ($p = 0.097$) or the graph theory technique ($p = 0.075$). The results from the graph theory, inter and intra observer techniques were not significantly different from each other ($p = 0.900$), achieving MSPE of 1.75, 1.78 and 1.32 respectively. One thing that is worth noting that unlike when looking at DSC where the inter and intra observer differences were very similar the intra observer is lower than the inter observer difference, while the difference is not significant ($p = 0.898$ for posterior surface) it is a possible indication that there is

some observer bias in the manual segmentation which may show up in a test with a larger number of images.

Table 21: 95% Hausdorff distance comparison between four automatic segmentation techniques and three manual segmentations

Segmentation Method	Mean	Standard Deviation
Threshold	25.39	8.89
CVWSiw	10.2	4.8
Graph Cut	6.73	2.65
Graph Theory	4.10	1.61
Inter Observer	3.58	0.95
Intra Observer	2.99	1.02

Table 21 above shows the results of 95% Hausdorff distance comparing the four different automated methods with manual segmentation results. A similar statistical analysis was carried out on the 95% Hausdorff distance. The results of this test were inconclusive. The only significant difference found was between the threshold method, with a 95% HD of 25.39, and all the other techniques ($p < 0.001$ for all). The rest of the techniques achieved values between 2.99 and 10.2. The different values of the mean 95% HD values were arranged in the same pattern seen previously. The largest value was found using the CVSWiw technique, the best of the automatic techniques was the graph theory method achieving a 95% HD value of 4.10. The differences that can be observed between the techniques are not statistically significant though. All of the methods have a reasonably large standard deviation (4.8 for CVWSiw, 1.61 for graph theory) reducing the significance of any differences. This is likely to be due to the fact that 95% HD looks only at the worst elements of the segmentation producing a wider variation in the results.

As well as accuracy another important measure of the suitability of a technique is the time it takes for segmentation to occur. The time taken for the different segmentation techniques is shown in Table 22 below. The four automatic techniques are compared. Two versions of the graph theory technique are shown, one using C++ and one using Matlab.

The fastest technique is the graph theory using C++ taking 0.288 seconds to segment an image. The next fastest methods are the threshold and graph cut segmentation techniques taking under 5 seconds to segment an image. Graph theory in Matlab was the next fastest technique taking just over 16 seconds to segment an image. The slowest technique was level set technique this took over

500 times longer than the graph theory technique. Time taken for manual segmentation was not measured but it would be expected to be around 15 minutes per image.

It is possible that further reductions in speed could be made in other methods that were not implemented solely C++. The Graph Cut method used both C++ and Matlab so while some reduction in speed could be achieved by sole use of C++ this is unlikely to be as dramatic as speed reductions achieved converting graph theory method from Matlab to C++.

Table 22: Comparison of time taken for four different automatic segmentation techniques

Method	Time (s)
Threshold	4.53
CVWSiw	150.7
Graph Cut	3.53
Graph Theory (using Matlab)	16.1
Graph Theory (using C++)	0.288

Overall looking at the different measures for accuracy similar results were achieved with all measures. All the newly developed segmentation techniques perform better than the previously used threshold technique. The graph cut and level set techniques have similar accuracy; however graph cut technique performs better than the level set technique when using the DSC technique. The graph theory technique is most accurate and doesn't show any significant differences with manual segmentation using any of the comparison techniques. The graph theory technique was the fastest technique. The superior performance in terms of both accuracy and speed of the graph theory based technique makes it the best technique.


5.6 3D Level Set

The 3D level set segmentation framework was tested on a synthetic image and a single real 3D AS-OCT volume. During the tests the constants determining the strength of different components of energy were empirically chosen for the best results. The values for the different constants used are $\lambda_1 = 0.2$, $\lambda_2^1 = \lambda_2^2 = 1$, $\lambda_3 = 0.8$ and $\lambda_4 = 0.1$.


A synthetic volume data was built using two ellipsoids with different radii to model the cornea. In each section some regions were deleted deliberately to simulate the OCT data where some regions are missing due to poor signal to noise ratio.

Speckle noise was also added to these images since speckle noise is present in OCT data. Figure 53(a) shows the synthetic data in montage form of all 32 sections. Figure 53(b) shows the segmentation results. It can be seen that the new segmentation framework is capable of recovering the artificial gaps we have put in our 'cornea' to model areas of lower signal found in real data.

The program was also tested on a single 3D OCT image of the human cornea, show in Figure 54(a). The image was taken from a healthy normal eye using a custom built spectral domain OCT machine. There are 32 cross-sectional scans of the cornea were taken with the scan being rotated 11.25 degrees between each image. Figure 54(b) shows the result of the segmentation. It can be seen from this that our program can segment the cornea.




(a)




(b)

Figure 53: A synthetic image and its segmentation result. (a) Synthetic data with 32 images arranged in a radial pattern; (b) Segmentation result of the image in (a).



(a)



(b)

Figure 54: Illustration of segmentation of a 3D AS-OCT image. (a) A 3D OCT image of the human cornea where 32 images shown here are arranged in a radial pattern to give the 3D image of the cornea; (b) Segmented cornea of the image (a).

While this work was able to achieve reasonable segmentation results it suffered from two major problems, the long time taken for the segmentation and the lack of a registration step. The program had a very large requirement for memory and so could only successfully segment smaller images. The results shown are on down sampled images. It also suffered from being too slow. The time taken to segment images of a single eye was over two hours which is not practical for medical usage. Registration of the images forming the volume was not considered by this method

which will lead to inaccuracies if the images are not perfectly aligned, a non uniform height of the corneal apex can be seen in Figure 54(a) indicating that a registration step is needed. Due to these problems work on 3D level set was abandoned and no numerical validation carried out.

5.7 3D Graph Cut

In order to validate the graph cut technique a repeatability test was carried out over a series of 17 eyes. For each eye two sets of images were acquired. The segmentation technique was carried out for the two sets of images. The results of the two segmentations were compared. Since we are imaging the same eye twice then the results should be the same. The comparison techniques used are described in section 1.6.2.

Since the segmentation techniques produce results with arbitrary coordinates the segmented images were aligned via their minimum point in the z direction prior to comparison.

Table 23 below shows the results of repeatability testing of this technique. Results are shown for three variations. The first is when all the images were cropped to the size of the smallest image and the other two when a progressively larger section of the image was also included. The results are best when only a small area is included; a mean difference of 2.10 pixels is found then compared to 3.62 for the image 150 pixels wider. The difference is probably due to the superior image quality near the centre of the images. There is also a smaller distance between points from different slices near the centre which will lead to better results in this region.

Table 23: Results of repeatability testing of 3D graph cut segmentation technique. Comparison of three variations of the alignment technique.

Method	Mean difference between surfaces		Repeatability of thickness	
	Mean	Standard deviation	Mean	Standard deviation
Graph Cut 3D segmentation	2.10	0.68	0.095	0.131
Graph Cut 3D segmentation 50 pixels wider	2.45	1.34	0.108	0.119

Graph Cut 3D segmentation 150 pixels wider	3.62	3.12	0.110	0.075
--	------	------	-------	-------

Table 24 below shows the time taken for each of three segmentation techniques. There is an increase in time taken for the methods that use a larger area of the images for the segmentation; the method covering the smallest area took 229 seconds compared to 235 seconds for method of larger area. This is unsurprising since we would expect it to take longer to segment larger images. The increase is fairly small, an increase in time taken of only 2.5% when an extra 150 pixels is added to each image. The reason the increase is so small is due to the fact that the initial registration steps are the same however the image is subsequently aligned. This registration process takes around 80% of the total time. The standard deviation is larger than the difference between the different techniques. This is due to a difference in the sizes of the different images that comes from the different locations where the eye lid had been previously cropped off.

Table 24: Time taken for segmentation for three variations of 3D graph cut segmentation technique. Results are taken over a set of 17 eyes.

Method	Time taken for segmentation (s)	
	Mean	Standard deviation
Graph Cut 3D segmentation	229	15
Graph Cut 3D segmentation 50 pixels wider	234	14
Graph Cut 3D segmentation 150 pixels wider	235	13

The results of the segmentation program show that it can achieve reasonable repeatability over a set of 17 eyes. The comparison of the different methods shows that the results are better towards the centre of the image and so are improved when a smaller area is used. The increase in time taken when using a larger area is small indicating that reducing the size of the images after the alignment step will not result in a significant increase in speed of segmentation. The decision on where to crop the images before segmentation must be made based on an assessment of

how useful the segmentation is in regions where the data points have a large space between them.

5.8 3D Graph Theory

The 3D graph theory work focussed on attempts to align the images correctly. The work here will be split into two sections. This first will cover a variety of different registration techniques that didn't use GA. The second section focuses on genetic algorithms and investigates a number of different energy functions that can be minimised.

5.8.1 Non GA Registration Techniques

All 11 registration methods were tested on the same data set. This consisted of two sets of 16 2D images from 17 eyes, giving a total of 554 images. The results from the alignment for each of the different sets were compared to each other. Two measures were used, mean difference between the surfaces and repeatability of thickness measurement across the surface.

Table 25: Repeatability comparison of 8 different alignment techniques over a set of 17 eyes. Mean difference measures the mean difference between the anterior surfaces of two different set of images of the same eye. Repeatability is the coefficient of repeatability of the thickness of the cornea.

Method	Mean Difference	Repeatability
No Reg	51.2	0.100
ICP	51.6	0.111
CP, straight line	14.1	0.073
CP, Curved	19.1	0.089
CP, ellipse	15.7	0.082
Pupil, straight line	22.4	0.195
Pupil, curved	26.5	0.196
Pupil, Ellipse	29.5	0.177

Table 25 above shows the results of the different alignment techniques. The best method, corner points with straight line achieved a mean difference of 14.1 pixels. All of the methods developed achieved better results than when no registration step was used. Using the ICP technique to align images showed no significant change in accuracy of alignment compared to using no registration, 51.6 compared to 51.2 pixels. This is probably due to the fact ICP is attempting to register the images to each other which is not the same as aligning all the images.

Techniques using corner points for the registration performed better than similar methods using the pupil endpoints, 14.1 pixels mean difference rather than 22.4. This is probably due to the detection of the corner points of the cornea being more reliable than the detection of the edges of the pupil. When looking at the different methods for using the points to align the images the most successful method is the straight line method. This holds true when using either corner points or pupil endpoints. When looking at corner points the ellipse method does better than the curved method for pupil endpoint method this is the other way round. The discrepancy is probably due to difficulties in correctly identifying the points. The straight line may work better since it provides a better estimate of the shape of the cornea than the curved line we were using. It is also less affected by small errors in positions of points so will perform better if there are some errors in locating the points.

Looking at the repeatability values the same results are not found. Here the best results are found using the corner points with straight line method to align the images. This is followed by the other corner points line methods. The differences between the different methods are much smaller when looking at the repeatability measure than the mean difference measure. These two measures are looking at different measurements which may explain some of the differences. One is measuring the mean difference between two anterior surfaces the other compares thickness measurements. The difference could be due to the fact the thickness doesn't alter as much across each image so small displacements don't show up as obviously.

In order to test the significance of the repeatability results an ANOVA test was carried out. This found that there was no significant difference between the results of different methods ($p = 0.345$). While the lack of significant difference may be partly due to the small sample size of the study it also indicates that the repeatability measure may not be a good way to distinguish between the methods developed.

When the ANOVA test was carried out on the mean difference data some significant differences between the methods were found. There was a significant difference between all the developed methods and the ICP and no registration methods ($p < 0.05$ for all).

5.8.2 Genetic Algorithm Energy Function Variations

Five different energy functions for use with GA were tested. The five different methods tested were circles in r , circles in r and z , ellipse fitting, ellipse fitting and

circles, and the model based method. The distance between two surfaces achieved by repeating method is used to evaluate the success of the method. Since the same eye has been imaged twice this distance should be zero. There will in reality be differences due to 3 different factors, the constancy of the OCT images (these could be caused by movement artefacts, limits to resolution etc.), the repeatability of the segmentation (some images may be slightly incorrectly segmented) and finally the errors may come from the registration program.

As well as measuring distance between the surfaces, the repeatability of thickness measurements was also calculated. The results of alignment are shown in Table 26.

Table 26: Comparison of performance of different GA models for alignment of 3D images.

Method	Mean difference between surfaces		Repeatability of thickness	
	Mean	Standard deviation	Mean	Standard deviation
Circles in r	23.12	12.4	0.0872	0.0430
Circles in r and z	18.45	12.4	0.0812	0.0372
Ellipsoid fitting	49.75	10.1	0.187	0.0572
Ellipsoid fitting and circles in r and z	20.18	15.9	0.102	0.0470
Model fitting	1.64	0.69	0.0877	0.0418
Best none GA method	14.1	3.2	0.073	0.035

Looking at the results for mean difference we can see that the first four of the energy functions for use with GA make the situation worse rather than better. Better results can be achieved by using only the initialisation step, mean difference 14.1, than these methods using a GA, mean difference between 18.45 and 49.75. The method involving the use of the Zernike polynomial model produced a substantial reduction in the mean difference achieving a mean difference of 1.64 pixels.

The results of the repeatability test for thickness were not as useful for distinguishing between the different models. Here the best results were achieved

by the non GA method; it also indicated that the circle fitting GA method was better than the model fitting technique though the difference between them is not significant. It is unclear why the two different techniques produce different results though one common feature of both evaluation techniques was that the ellipsoid based methods performed worst.

5.8.3 Genetic Algorithm Range Variations

In addition to testing different energy functions five different ways to limit the range of the values for the solution were also tested. These were three of the non GA methods, no limit and limit applied by fitting ellipses to four subsets of images.

Table 27 shows the results of the alignment.

Table 27: Comparison of different methods of limiting the range of possible solutions when using GA. Abbreviations refer to methods shown in Table 25 above which have been used here to limit the range of possible solutions for the GA.

Method	Mean Difference	Repeatability
CPSL limit	4.47	0.079
CPC limit	6.39	0.100
CPSL limit xy only	1.77	0.085
No limit	2.52	0.100
Four ellipsoids limit	60.2	n/a

The method with the limit from straight line in x and y directions only was the best performing technique achieving a mean difference of 1.77. When the z value was limited by the initialisation this produced worse results. This shows that the initialisation techniques are not able to accurately find the shift in the z direction. The performance was better when no limit was applied after the initial step than when the initial step limited the range of z values.

Looking at the repeatability results the best performing technique was the corner points and straight line method of setting limits with limits on all dimensions. This does not agree with what was found when looking at the mean difference of anterior surface. This discrepancy between mean difference and thickness repeatability has been found looking at all the techniques and is likely due to the fact thickness varies less than height across the cornea so it is harder to detect misalignment looking at thickness as opposed to height.

The method using four intersection points of images to improve the alignment produced the worse results of all the methods, with a mean difference of 60.2,

indicating that this was not a good method for further testing. No thickness measurements were made for this technique due to its very poor results in aligning the anterior surface. No attempt was made at using this to align the posterior surface needed to make thickness measurements.

5.9 Comparison of 3D Segmentation Techniques

Three different segmentation techniques were attempted for 3D segmentation with there being a number of variations of each method. The different methods were level set segmentation, graph cut segmentation and graph theory segmentation with genetic algorithm alignment.

A full comparison was not carried out of the level set method since this was unable to segment full sized images due to requiring too much memory to run on computer. The large memory requirements and the slowness of the method meant it was rejected as a possible technique. A quantitative comparison was carried out to compare graph cut segmentation, graph theory segmentation with genetic algorithm alignment and graph theory segmentation with non GA alignment.

Table 28: Comparison of three different 3D segmentation techniques over a set of 17 eyes

Method	Mean difference between surfaces		Repeatability of thickness	
	Mean	Standard deviation	Mean	Standard deviation
Graph Cut	2.45	1.34	0.108	0.119
Graph Theory with GA	1.64	0.69	0.0877	0.0418
Graph Theory without GA	14.1	3.2	0.073	0.035

Table 28 shows the results of a comparison between the performances of three different 3D segmentation techniques. Looking at the results of the mean difference test it can be seen that the best results are achieved using the graph theory and GA technique, a mean difference of 1.64 pixels was found. This performs substantially better than using graph theory segmentation with identified points on the image to register the images, 1.64 pixel difference compared to 14.1 pixels. The graph cut technique performs almost as well as the graph theory and GA technique, achieving a mean difference of 2.45 pixels. The difference between the two methods is not significant ($p < 0.05$). Looking at the repeatability results a different pattern is seen.

Here the graph theory without GA performs best and graph cut technique performs worst. The difference between the two evaluation techniques is possibly because repeatability is a measure that looks at both anterior and posterior surfaces whereas mean difference only looks at the anterior surface. The non GA technique uses reference points on the posterior surface while the GA method uses only points on the anterior surface for alignment.

Table 29: Time taken for segmentation and alignment of 3D images. Results taken over set of 17 eyes.

Method	Time taken for segmentation and alignment (seconds)	
	Mean	Standard deviation
Graph Cut	233	14
Graph Theory with GA	176	15
Graph Theory without GA	148	9

Table 29 above shows how long segmentation and alignment took for the different methods. The fastest method is the graph theory without GA, taking 148 seconds. The slowest method is the graph cut technique, taking 233 seconds. The graph cut technique is probably slowest since it performs two segmentation steps an initial 2D segmentation which is used for registration then a full 3D segmentation. The other two methods only perform 2D segmentation followed by an alignment step. Since the GA method uses genetic algorithms to improve the results of the initial method it will always take longer than if the GA step doesn't take place. All the methods are able to complete segmentation and alignment in less than 4 minutes. This means they would be able to be easily used to give results during a patient visit without having to wait a long time for analysis of images to take place.

Repeatability tests were chosen for a measure of the success of the technique for a number of reasons. If a technique is to ever be useful then measurements of a particular object should be the same each time the object is imaged. A device that produces different measurements of the same object each time you measure it is not very useful. For the 2D segmentation analysis the results of the automatic segmentation were compared with manual segmentation. Comparison with manual observers is not practical for 3D segmentation. Manual registration of the images

would be a very time consuming process and is not one which is commonly carried out.

An alternative method for evaluating the data would be to image the eye using an alternative imaging modality and compare the results on this method with those achieved using a different technique. A suitable set of images, where a person had been imaged both with OCT and another more established technique for measuring the human cornea, was not available.

The repeatability of the thickness measurement has not been found to be a good metric for judging the effectiveness of the technique. The results do not agree with those found looking at a mean surface difference and don't show significant differences between the different methods.

Overall looking at the quality of the results and the time taken for segmentation for the different techniques the method that performs best is the graph theory with genetic algorithm method. This achieves the best results when looking at the mean difference comparison where it significantly outperforms the graph theory without GA. It is also faster than the graph cut method and achieves better repeatability than the graph cut method when looking at thickness measurements.

The values achieved by the GA model method were similar to the difference between the segmentation technique and manual segmentation. This indicates that it is not likely to be able to achieve further improvement to the alignment technique as there will always be an error present from segmentation of the images.

This is a reasonably quick method taking around 3 minutes to achieve 3D segmentation. It gives information of the boundaries of the anterior segment across the entire region that the images are able to cover.

6 Discussion and Conclusion

6.1 Discussion

A number of novel methods for the segmentation of anterior segment OCT images have been developed for this study. For 2D segmentation these include techniques using level set, graph cut and graph theory based methods. Analysis of the results of segmentation by these methods found the most successful technique was graph theory with shape. This method has been shown to achieve results with the same level of accuracy as expert manual segmentation, achieving a DSC value of 0.961 when compared to manual segmentation. The graph theory with shape method is a fast method, taking 0.288 seconds to segment a single image.

This new technique has been compared to a number of previously used segmentation techniques. It has been shown to perform significantly better than a previously used threshold method [29]. The threshold technique achieved a DSC value of 0.852. Another method that has been previously used is a graph theory based segmentation method. This differs from our method in that it does not include a shape term as part of the segmentation. The group that first published this technique were only looking at the central region of the cornea where the highest signal to noise ratio is found [97]. The results of the graph theory with shape technique developed here were compared to the same technique without the shape term. It was found that without the shape technique, the method was unable to accurately detect the posterior surface of the cornea: a posterior MSPE of 17.3 pixels was found compared to 1.49 pixels when using the shape term. This indicates the newly developed method is able to accurately segment a larger region of the cornea than previously developed versions of this method.

Tomey have developed a swept source OCT machine, CASIA (Tomey, Tokyo, Japan). Swept source OCT is a variety of SD OCT. This machine comes with algorithms that are able to automatically segment both the anterior and posterior surfaces across the entire anterior segment. They have used a proprietary method to carry out segmentation and there is no publically available information on their method. A comparison of our technique to the technique they have developed would be of interest. It has not been possible to source any images using this machine, so this has not been done. The development of accessible methods where the technique used is made public through publication is important for the use of segmentation in research purposes. Having a method that can be adapted to target specific parts of the anterior segment or eyes with specific diseases is useful.

Further advances in OCT technology will affect the work. Most of the images acquired for this work were taken using time domain OCT. Spectral domain OCT is a newer technique that is able to achieve higher speed and resolution images. Current commercially available SD OCT machines have a shorter scanning depth than can be achieved using time domain techniques. This means they are unable to image the entire anterior segment and can only produce images of small areas. They can be focussed on central regions or peripheral regions depending on what is being investigated. This makes them unsuited for use to acquire topography information for the entire cornea. Further advances in OCT technology are likely. There are a number of groups that have custom built OCT devices that are capable of imaging a larger area. It is likely that these devices will become more widely available in the not too distant future.

Spectral domain images are generally better quality than time domain images. This means that the segmentation method developed here should also work with images from any new spectral domain OCT machines that are developed. The fact programs were developed to work on time domain OCT images should not provide a drawback. It may offer an advantage in that it is able to segment images that contain regions of low signal to noise ratio. When medical imaging is carried out on a large scale it is likely that there will be images that have regions that do not have optimum image quality. Our program should still be able to segment these regions.

For 3D segmentation, the best technique developed used graph theory with shape to segment 2D images then used a novel alignment method to align images to produce a 3D surface. This technique used identification of three key points as an initial guide followed by the use of genetic algorithms to minimise the distance between the surface and a model surface of the cornea. This has been shown to achieve a mean difference of 1.7 pixels between two sets of images of the same eye.

There has been only very limited previous work looking at 3D segmentation of OCT image of the cornea. No quantitative validation of 3D segmentation has been carried out by other groups previously investigating the cornea. Eichel et al present a method for segmenting and aligning 3D SDOCT images [131]. Their study carries out validation tests for the 2D segmentation process but they only have data from 3 eyes for their 3D work so do not present any validation of the reliability of their technique. The work presented here is the first time validation of 3D alignment and segmentation of OCT images of the cornea has been carried out.

One of the limitations of the work was the relatively small number of images that have been used to evaluate the technique. The 3D method has only been carried out on a sample size of 17 eyes. Analysis of results over a larger data set would allow for more confidence to be had in the conclusions from the study and may also lead to further improvements to the model.

Combining these techniques creates a tool that can go from a set of OCT images of the human eye to a map showing the location of the anterior and posterior surfaces across the entire cornea and including part of the sclera. This tool will allow clinicians to use information acquired from OCT images of the human eye without the need for lengthy manual analysis of images. The maps that are created can be used as an input for patient specific biomechanical modelling of the cornea for the purpose of treatment, planning and disease monitoring. The speed of the segmentation is sufficiently fast to allow for analysis of results while a patient waits.

6.2 Conclusion

Three different novel techniques have been developed for the segmentation of 2D OCT images of the anterior segment. These are level set with shape, graph cut with shape and graph theory with shape. Analysis of the results of these methods found that the most successful technique was graph theory with shape. This method achieved similar performance to manual segmentation and showed superior performance when compared to previously used techniques. For 3D segmentation, each of these 2D techniques was extended to 3D. The level set and graph cut techniques could be extended to true 3D segmentation. The graph theory technique used “2.5D” segmentation where a series of 2D images are aligned to produce a 3D volume. Analysis of the results of these methods found that the most successful technique was graph theory segmentation with registration using genetic algorithms. This method showed good repeatability. There has been no previous work validating registration and segmentation of 3D OCT images of the cornea. When combined, a method has been developed that can take a series of 2D OCT images of the anterior segment of the human eye and produce a map of the anterior and posterior surfaces of the cornea.

6.3 Future Work

6.3.1 Comparison with Results from Other Imaging Modalities

In order to further develop the work discussed here, further validation of the techniques could be carried out. So far no comparison of our results with the results

that can be achieved using different imaging modalities has occurred. Given the strengths and weaknesses of different imaging technologies it is likely that a variety of imaging techniques will continue to be used to investigate the cornea. A comparison between our segmentation of OCT images and results that can be achieved through Scheimplug photography would be useful when trying to decide what imaging modality is more useful for a patient's specific problem.

6.3.2 Segmentation of SD OCT Images

If a different imaging system was used there would be a number of minor changes to the model needed. The main change that would need to be made would be adjusting coefficients chosen for various parameters. Specifically the creation of shape terms to guide the segmentation of the posterior surface of the cornea used coefficients to predict how thickness varied with distance from centre. Different sized images would require altering these coefficients. The optimisation of various other parameters may also need updating if a different imaging system was used.

6.3.3 Incorporation with Biomechanical Modelling

The techniques that have been developed here have had the aim of segmenting OCT images so they can be used to help produce patient specific biomechanical models of the human eye. 3D maps showing the location of the anterior and posterior surfaces of the cornea are produced by methods developed over the course of this PhD. A relatively simple extension to the work would be to incorporate these 3D maps into biomechanical modelling.

6.3.4 Improvements to the Segmentation Technique

The 2D segmentation technique developed here performs very well and has limited potential for further improvements to be made on the model. The main scope for further development of the technique lies in the 3D alignment techniques. The method developed here uses a model cornea produced from fitting a Zernike polynomial to a corneal surface. This model is then used to align the images in an appropriate manner. It is possible for improvements to be made to this model cornea. So far a relatively small data set has been used to test the method. A larger data set could be used to produce a more accurate model of the cornea. This could be used to guide the alignment of the images to produce a 3D shape.

7 References

- [1] D. Williams, Y. Zheng, F. Bao, and A. Elsheikh, "Automatic segmentation of anterior segment optical coherence tomography images," *Journal of Biomedical Optics*, vol. 18, pp. 056003-056003, 2013.
- [2] D. Williams, Y. Zheng, F. Bao, and A. Elsheikh, "Fast segmentation of anterior segment optical coherence tomography images using graph cut," *Eye and Vision*, vol. 2, p. 1, 2015.
- [3] A. S. Roy and W. J. Dupps Jr, "Patient-specific modeling of corneal refractive surgery outcomes and inverse estimation of elastic property changes," *Journal of biomechanical engineering*, vol. 133, p. 011002, 2011.
- [4] D. M. Maurice, "The structure and transparency of the cornea," *The Journal of physiology*, vol. 136, pp. 263-286, 1957.
- [5] A. D. D. John V. Forrester, Paul G. Mcmenamin, William R. Lee, *The eye basic sciences in practice*, 2nd ed.: Saunders, 2002.
- [6] N. I. o. H. National Eye Institute. *Eye diagram showing the macula and fovea (black and white)*. Available: http://www.nei.nih.gov/photo/eyean/images/NEA09_72.jpg
- [7] Y. Komai and T. Ushiki, "The three-dimensional organization of collagen fibrils in the human cornea and sclera," *Investigative Ophthalmology and Visual Science*, vol. 32, pp. 2244-2258, 1991.
- [8] M. Wang, *LASIK vision correction - an overview of refractive errors and their treatment with the excimer laser*: Med World Publishing, 2000.
- [9] L. R. Dice, "Measures of the Amount of Ecologic Association Between Species," *Ecology*, vol. 26, pp. 297-302, 1945/07/01 1945.
- [10] S. Klein, U. A. Van Der Heide, I. M. Lips, M. Van Vulpen, M. Staring, and J. P. W. Pluim, "Automatic segmentation of the prostate in 3D MR images by atlas matching using localized mutual information," *Medical Physics*, vol. 35, pp. 1407-1417, 2008.
- [11] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, "Comparing images using the Hausdorff distance," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, pp. 850-863, 1993.
- [12] N. Archip, O. Clatz, S. Whalen, D. Kacher, A. Fedorov, A. Kot, *et al.*, "Non-rigid alignment of pre-operative MRI, fMRI, and DT-MRI with intra-operative MRI for enhanced visualization and navigation in image-guided neurosurgery," *NeuroImage*, vol. 35, pp. 609-624, 2007.
- [13] M. J. Doughty and M. L. Zaman, "Human corneal thickness and its impact on intraocular pressure measures: A review and meta-analysis approach," *Survey of Ophthalmology*, vol. 44, pp. 367-408, 2000.
- [14] T. Olsen, C. B. Nielsen, and N. Ehlers, "On the optical measurement of corneal thickness. I. Optical principle and sources of error," *Acta Ophthalmologica*, vol. 58, pp. 760-766, 1980.
- [15] S. Patel and R. W. W. Stevenson, "Clinical evaluation of a portable ultrasonic and a standard optical pachometer," *Optometry and Vision Science*, vol. 71, pp. 43-46, 1994.
- [16] T. Olsen and N. Ehlers, "The thickness of the human cornea as determined by a specular method," *Acta Ophthalmologica*, vol. 62, pp. 859-871, 1984.
- [17] R. Bovelie, S. C. Kaufman, H. W. Thompson, and H. Hamano, "Corneal thickness measurements with the Topcon SP-2000P specular microscope and an ultrasound pachymeter," *Archives of Ophthalmology*, vol. 117, pp. 868-870, 1999.
- [18] S. Suzuki, T. Oshika, K. Oki, I. Sakabe, A. Iwase, S. Amano, *et al.*, "Corneal thickness measurements: scanning-slit corneal topography and noncontact specular microscopy versus ultrasonic pachymetry," *Journal of Cataract and Refractive Surgery*, vol. 29, pp. 1313-1318, 2003.

- [19] Z. Liu, A. J. Huang, and S. C. Pflugfelder, "Evaluation of corneal thickness and topography in normal eyes using the Orbscan corneal topography system," *British Journal of Ophthalmology*, vol. 83, pp. 774-778, 1999.
- [20] M. M. Marsich and M. A. Bullimore, "The repeatability of corneal thickness measures," *Cornea*, vol. 19, pp. 792-795, 2000.
- [21] M. A. Fakhry, A. Artola, J. I. Belda, M. J. Ayala, and J. L. Alió, "Comparison of corneal pachymetry using ultrasound and Orbscan II," *Journal of Cataract and Refractive Surgery*, vol. 28, pp. 248-252, 2002.
- [22] J. González-Pérez, J. M. González-Méijome, M. T. Rodríguez Ares, and M. A. Parafita, "Central corneal thickness measured with three optical devices and ultrasound pachometry," *Eye and Contact Lens*, vol. 37, pp. 66-70, 2011.
- [23] A. Gordon, E. A. Boggess, and J. F. Molinari, "Variability of ultrasonic pachometry," *Optometry and Vision Science*, vol. 67, pp. 162-165, 1990.
- [24] H. Ishikawa and J. S. Schuman, "Anterior segment imaging: Ultrasound biomicroscopy," *Ophthalmology Clinics of North America*, vol. 17, pp. 7-20, 2004.
- [25] W. Nolan, "Anterior segment imaging: Ultrasound biomicroscopy and anterior segment optical coherence tomography," *Current Opinion in Ophthalmology*, vol. 19, pp. 115-121, 2008.
- [26] A. Wegener and H. Laser-junga, "Photography of the anterior eye segment according to Scheimpflug's principle: Options and limitations - A review," *Clinical and Experimental Ophthalmology*, vol. 37, pp. 144-154, 2009.
- [27] D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, *et al.*, "Optical coherence tomography," *Science*, vol. 254, pp. 1178-1181, 1991.
- [28] J. A. Izatt, M. R. Hee, E. A. Swanson, C. P. Lin, D. Huang, J. S. Schuman, *et al.*, "Micrometer-scale resolution imaging of the anterior eye in vivo with optical coherence tomography," *Archives of Ophthalmology*, vol. 112, pp. 1584-1589, 1994.
- [29] L. Ge, M. Shen, A. Tao, J. Wang, G. Dou, and F. Lu, "Automatic segmentation of the central epithelium imaged with three optical coherence tomography devices," *Eye & Contact Lens: Science & Clinical Practice*, 2012.
- [30] W. Drexler, U. Morgner, R. K. Ghanta, F. X. Kärtner, J. S. Schuman, and J. G. Fujimoto, "Ultrahigh-resolution ophthalmic optical coherence tomography," *Nature Medicine*, vol. 7, pp. 502-506, 2001.
- [31] S. Sin and T. L. Simpson, "The repeatability of corneal and corneal epithelial thickness measurements using optical coherence tomography," *Optometry and Vision Science*, vol. 83, pp. 360-365, 2006.
- [32] S. A. Boppart, G. J. Tearney, B. E. Bouma, J. F. Southern, M. E. Brezinski, and J. G. Fujimoto, "Noninvasive assessment of the developing Xenopus cardiovascular system using optical coherence tomography," *Proceedings of the National Academy of Sciences*, vol. 94, pp. 4256-4261, 1997.
- [33] M. E. J. van Velthoven, D. J. Faber, F. D. Verbraak, T. G. van Leeuwen, and M. D. de Smet, "Recent developments in optical coherence tomography for imaging the retina," *Progress in Retinal and Eye Research*, vol. 26, pp. 57-77, 2007.
- [34] W. F. Cheong, S. A. Prahl, and A. J. Welch, "A review of the optical properties of biological tissues," *IEEE Journal of Quantum Electronics*, vol. 26, pp. 2166-2185, 1990.
- [35] B. J. Kaluzy, J. J. Kałuzny, A. Szkulmowska, I. Gorczyńska, M. Szkulmowski, T. Bajraszewski, *et al.*, "Spectral optical coherence tomography: A novel technique for cornea imaging," *Cornea*, vol. 25, pp. 960-965, 2006.
- [36] E. Wylegała, S. Teper, A. K. Nowińska, M. Milka, and D. Dobrowolski, "Anterior segment imaging: fourier-domain optical coherence tomography versus time-

- domain optical coherence tomography," *Journal of Cataract and Refractive Surgery*, vol. 35, pp. 1410-1414, 2009.
- [37] (2012). *Scopus Database*. Available: <http://www.scopus.com>
- [38] J. L. B. Ramos, Y. Li, and D. Huang, "Clinical and research applications of anterior segment optical coherence tomography - A review," *Clinical and Experimental Ophthalmology*, vol. 37, pp. 81-89, 2009.
- [39] M. Doors, T. T. J. M. Berendschot, J. de Brabander, C. A. B. Webers, and R. M. M. A. Nuijts, "Value of optical coherence tomography for anterior segment surgery," *Journal of Cataract and Refractive Surgery*, vol. 36, pp. 1213-1229, 2010.
- [40] C. Du, M. Shen, M. Li, D. Zhu, M. R. Wang, and J. Wang, "Anterior segment biometry during accommodation imaged with ultralong scan depth optical coherence tomography," *Ophthalmology*, vol. 119, pp. 2479-2485, 2012.
- [41] H. Quigley and A. T. Broman, "The number of people with glaucoma worldwide in 2010 and 2020," *British Journal of Ophthalmology*, vol. 90, pp. 262-267, 2006.
- [42] R. Sihota, P. Vashisht, A. Sharma, S. Chakraborty, V. Gupta, and R. M. Pandey, "Anterior segment optical coherence tomography characteristics in an Asian population," *Journal of Glaucoma*, vol. 21, pp. 180-185, 2012.
- [43] J. Tian, P. Marziliano, M. Baskaran, H. T. Wong, and T. Aung, "Automatic anterior chamber angle assessment for HD-OCT images," *IEEE Transactions on Biomedical Engineering*, vol. 58, pp. 3242-3249, 2011.
- [44] S. Liu, M. Yu, C. Ye, D. S. C. Lam, and C. K. S. Leung, "Anterior chamber angle imaging with swept-source optical coherence tomography: An investigation on variability of angle measurement," *Investigative Ophthalmology and Visual Science*, vol. 52, pp. 8598-8603, 2011.
- [45] M. Q. Salomão, A. Esposito, and W. J. Dupps Jr, "Advances in anterior segment imaging and analysis," *Current Opinion in Ophthalmology*, vol. 20, pp. 324-332, 2009.
- [46] B. Qin, M. Tang, Y. Li, X. Zhang, R. Chu, and D. Huang, "Anterior segment dimensions in Asian and Caucasian eyes measured by optical coherence tomography," *Ophthalmic Surgery Lasers and Imaging*, vol. 43, pp. 135-142, 2012.
- [47] A. Konstantopoulos, J. Kuo, D. Anderson, and P. Hossain, "Assessment of the use of anterior segment optical coherence tomography in microbial keratitis," *American Journal of Ophthalmology*, vol. 146, pp. 534-542.e2, 2008.
- [48] C. J. Pavlin, L. M. Vásquez, R. Lee, E. R. Simpson, and I. I. K. Ahmed, "Anterior segment optical coherence tomography and ultrasound biomicroscopy in the imaging of anterior segment tumors," *American Journal of Ophthalmology*, vol. 147, pp. 214-219.e2, 2009.
- [49] G. Prakash, D. Ashokumar, S. Jacob, K. Satish Kumar, A. Agarwal, and A. Agarwal, "Anterior segment optical coherence tomography-aided diagnosis and primary posterior chamber intraocular lens implantation with fibrin glue in traumatic phacocoele with scleral perforation," *Journal of Cataract and Refractive Surgery*, vol. 35, pp. 782-784, 2009.
- [50] Y. Li, M. V. Netto, R. Shekhar, R. R. Krueger, and D. Huang, "A longitudinal study of LASIK flap and stromal thickness with high-speed optical coherence tomography," *Ophthalmology*, vol. 114, pp. 1124-1132.e1, 2007.
- [51] R. C. Hall, F. K. Mohamed, H. M. Htoon, D. T. Tan, and J. S. Mehta, "Laser in situ keratomileusis flap measurements: Comparison between observers and between spectral-domain and time-domain anterior segment optical coherence tomography," *Journal of Cataract and Refractive Surgery*, vol. 37, pp. 544-551, 2011.

- [52] M. M. Lai, M. Tang, E. M. M. Andrade, Y. Li, R. N. Khurana, J. C. Song, *et al.*, "Optical coherence tomography to assess intrastromal corneal ring segment depth in keratoconic eyes," *Journal of Cataract and Refractive Surgery*, vol. 32, pp. 1860-1865, 2006.
- [53] M. Doors, N. G. Tahzib, F. A. Eggink, T. T. J. M. Berendschot, C. A. B. Webers, and R. M. M. A. Nuijts, "Use of anterior segment optical coherence tomography to study corneal changes after collagen cross-linking," *American Journal of Ophthalmology*, vol. 148, pp. 844-851.e2, 2009.
- [54] M. Jancevski and C. S. Foster, "Anterior segment optical coherence tomography," *Seminars in Ophthalmology*, vol. 25, pp. 317-323, 2010.
- [55] J. P. S. Garcia Jr, J. De La Cruz, R. B. Rosen, and D. F. Buxton, "Imaging implanted keratoprotheses with anterior-segment optical coherence tomography and ultrasound biomicroscopy," *Cornea*, vol. 27, pp. 180-188, 2008.
- [56] L. S. Lim, H. T. Aung, T. Aung, and D. T. H. Tan, "Corneal imaging with anterior segment optical coherence tomography for lamellar keratoplasty procedures," *American Journal of Ophthalmology*, vol. 145, pp. 81-90, 2008.
- [57] C. M. Prospero Ponce, K. M. Rocha, S. D. Smith, and R. R. Krueger, "Central and peripheral corneal thickness measured with optical coherence tomography, Scheimpflug imaging, and ultrasound pachymetry in normal, keratoconus-suspect, and post-laser in situ keratomileusis eyes," *Journal of Cataract and Refractive Surgery*, vol. 35, pp. 1055-1062, 2009.
- [58] J. Huang, X. Ding, G. Savini, C. Pan, Y. Feng, D. Cheng, *et al.*, "A comparison between scheimpflug imaging and optical coherence tomography in measuring corneal thickness," *Ophthalmology*, vol. 120, pp. 1951-1958, 2013.
- [59] D. Z. Reinstein, M. Gobbe, and T. J. Archer, "Anterior segment biometry: A study and review of resolution and repeatability data," *Journal of Refractive Surgery*, vol. 28, pp. 509-520, 2012.
- [60] R. E. W. Rafael C. Gonzalez, *Digital image processing*, 3rd Edition ed.: Person Education, 2008.
- [61] N. Otsu, "Threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-9, pp. 62-66, 1979.
- [62] J. L. Semmlow, *Biosignal and medical image processing*, Second Edition ed.: CRC Press, 2009.
- [63] M. Shen, L. Cui, M. Li, D. Zhu, M. R. Wang, and J. Wang, "Extended scan depth optical coherence tomography for evaluating ocular surface shape," *Journal of Biomedical Optics*, vol. 16, 2011.
- [64] L. M. Sakata, R. Lavanya, D. S. Friedman, H. T. Aung, H. Gao, R. S. Kumar, *et al.*, "Comparison of gonioscopy and anterior segment ocular coherence tomography in detecting angle closure in different quadrants of the anterior chamber angle," *Ophthalmology*, vol. 115, pp. 769-774, 2008.
- [65] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *International Journal of Computer Vision*, vol. 1, pp. 321-331, 1988.
- [66] R. Malladi, J. A. Sethian, and B. C. Vemuri, "Shape modeling with front propagation: a level set approach," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, pp. 158-175, 1995.
- [67] S. Osher and R. P. Fedkiw, "Level set methods: an overview and some recent results," *Journal of Computational Physics*, vol. 169, pp. 463-502, 2001.
- [68] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *International Journal of Computer Vision*, vol. 22, pp. 61-79, 1997.

- [69] D. C. Fernández, "Delineating fluid-filled region boundaries in optical coherence tomography images of the retina," *IEEE Transactions on Medical Imaging*, vol. 24, pp. 939-945, 2005.
- [70] T. F. Chan and L. A. Vese, "Active contours without edges," *IEEE Transactions on Image Processing*, vol. 10, pp. 266-277, 2001.
- [71] M. Rousson and N. Paragios, "Prior knowledge, level set representations & visual grouping," *International Journal of Computer Vision*, vol. 76, pp. 231-243, 2008.
- [72] P. H. Lim, U. Bagci, and L. Bai, *A new prior shape model for level set segmentation* vol. 7042 LNCS, 2011.
- [73] A. Yazdanpanah, G. Hamarneh, B. Smith, and M. Sarunic, *Intra-retinal layer segmentation in optical coherence tomography using an active contour approach* vol. 5762 LNCS, 2009.
- [74] D. Cremers, M. Rousson, and R. Deriche, "A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape," *International Journal of Computer Vision*, vol. 72, pp. 195-215, 2007.
- [75] S. Chen and R. J. Radke, "Level set segmentation with both shape and intensity priors," 2009, pp. 763-770.
- [76] A. Tsai, A. Yezzi Jr, W. Wells, C. Tempany, D. Tucker, A. Fan, *et al.*, "A shape-based approach to the segmentation of medical imagery using level sets," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 137-154, 2003.
- [77] X. Bresson, P. Vanderghelynst, and J. P. Thiran, "A variational model for object segmentation using boundary information and shape prior driven by the Mumford-Shah functional," *International Journal of Computer Vision*, vol. 68, pp. 145-162, 2006.
- [78] M. E. Leventon, W. E. L. Grimson, and O. Faugeras, "Statistical shape influence in geodesic active contours," in *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, 2000, pp. 316-323 vol.1.
- [79] A. Yazdanpanah, G. Hamarneh, B. R. Smith, and M. V. Sarunic, "Segmentation of intra-retinal layers from optical coherence tomography images using an active contour approach," *IEEE Transactions on Medical Imaging*, vol. 30, pp. 484-496, 2011.
- [80] C. Li, C. Xu, C. Gui, and M. D. Fox, "Level set evolution without re-initialization: a new variational formulation," 2005, pp. 430- 436 vol. 1-430- 436 vol. 1.
- [81] C. Pluempitwiriwawej, J. M. F. Moura, Y. J. L. Wu, and C. Ho, "STACS: new active contour scheme for cardiac MR image segmentation," *IEEE Transactions on Medical Imaging*, vol. 24, pp. 593-603, 2005.
- [82] Z. X., C. W., and P. Q., "Efficiently solving the piecewise constant mumford-shah model using graph cuts " *Technical report, Dept. of Computer Science, Zhejiang University, P.R. China,,* 2006.
- [83] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, pp. 1222-1239, 2001.
- [84] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, pp. 1124-1137, 2004.
- [85] V. Kolmogorov and R. Zabih, "What Energy Functions Can Be Minimized via Graph Cuts?," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, pp. 147-159, 2004.
- [86] O. Daněš and P. Matula, "An improved Riemannian metric approximation for graph cuts," in *Lecture Notes in Computer Science (including subseries Lecture Notes in*

- Artificial Intelligence and Lecture Notes in Bioinformatics*) vol. 6607 LNCS, ed, 2011, pp. 71-82.
- [87] N. El-Zehiry, S. Xu, P. Sahoo, and A. Elmaghraby, "Graph cut optimization for the mumford-shah model," in *7th IASTED International Conference on Visualization, Imaging, and Image Processing*, Palma de Mallorca; Spain, 2007, pp. 182-187.
- [88] G. Slabaugh and G. Unal, "Graph cuts segmentation using an elliptical shape prior," in *IEEE International Conference on Image Processing*, Genova; Italy, 2005, pp. 1222-1225.
- [89] O. Daněk, P. Matula, M. Maška, and M. Kozubek, "Smooth Chan-Vese segmentation via graph cuts," *Pattern Recognition Letters*, vol. 33, pp. 1405-1410, 2012.
- [90] A. M. Ali, A. A. Farag, and A. S. El-Baz, "Graph cuts framework for kidney segmentation with prior shape constraints," vol. 4791 LNCS, ed, 2007, pp. 384-392.
- [91] M. K. Garvin, M. D. Abramoff, R. Kardon, S. R. Russell, X. Wu, and M. Sonka, "Intraretinal layer segmentation of macular optical coherence tomography images using optimal 3-D graph search," *IEEE Transactions on Medical Imaging*, vol. 27, pp. 1495-1505, 2008.
- [92] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269-271, 1959.
- [93] V. H. M. Sonka, R. Boyle, *Image Processing, Analysis and Machine vision*: Thomson Learning, 2008.
- [94] S. J. Chiu, X. T. Li, P. Nicholas, C. A. Toth, J. A. Izatt, and S. Farsiu, "Automatic segmentation of seven retinal layers in SDOCT images congruent with expert manual segmentation," *Optics Express*, vol. 18, pp. 19413-19428, 2010.
- [95] Q. Yang, C. A. Reisman, Z. Wang, Y. Fukuma, M. Hangai, N. Yoshimura, *et al.*, "Automated layer segmentation of macular OCT images using dual-scale gradient information," *Optics Express*, vol. 18, pp. 21293-21307, 2010.
- [96] S. J. Chiu, C. A. Toth, C. B. Rickman, J. A. Izatt, and S. Farsiu, "Automatic segmentation of closed-contour features in ophthalmic images using graph theory and dynamic programming," *Biomedical Optics Express*, vol. 3, pp. 1127-1140, 2012.
- [97] F. LaRocca, S. J. Chiu, R. P. McNabb, A. N. Kuo, J. A. Izatt, and S. Farsiu, "Robust automatic segmentation of corneal layer boundaries in SDOCT images using graph theory and dynamic programming," *Biomedical Optics Express*, vol. 2, pp. 1524-1538, 2011.
- [98] M. Baroni, P. Fortunato, and A. La Torre, "Towards quantitative analysis of retinal features in optical coherence tomography," *Medical Engineering & Physics*, vol. 29, pp. 432-441, 2007.
- [99] M. Wojtkowski, V. Srinivasan, J. G. Fujimoto, T. Ko, J. S. Schuman, A. Kowalczyk, *et al.*, "Three-dimensional retinal imaging with high-speed ultrahigh-resolution optical coherence tomography," *Ophthalmology*, vol. 112, pp. 1734-1746, 2005.
- [100] S. Jiao, R. Knighton, X. Huang, G. Gregori, and C. A. Puliafito, "Simultaneous acquisition of sectional and fundus ophthalmic images with spectral-domain optical coherence tomography," *Optics Express*, vol. 13, pp. 444-452, 2005.
- [101] J. Yang and J. S. Duncan, "3D image segmentation of deformable objects with joint shape-intensity prior models using level sets," *Medical Image Analysis*, vol. 8, pp. 285-294, 2004.
- [102] M. Haeker, M. Sonka, R. Kardon, V. A. Shah, X. Wu, and M. D. Abramoff, "Automated segmentation of intraretinal layers from macular optical coherence tomography images," in *Progress in Biomedical Optics and Imaging - Proceedings of SPIE*, 2007.

- [103] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 239-256, 1992.
- [104] S. Du, N. Zheng, S. Ying, and J. Liu, "Affine iterative closest point algorithm for point set registration," *Pattern Recognition Letters*, vol. 31, pp. 791-799, 2010.
- [105] T. Pajdla and L. V. Gool, "Matching of 3-D curves using semi-differential invariants," in *IEEE International Conference on Computer Vision*, 1995, pp. 390-395.
- [106] C. S. Lee, S. M. Guo, and C. Y. Hsu, "Genetic-based fuzzy image filter and its application to image processing," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 35, pp. 694-711, 2005.
- [107] C. Munteanu and A. Rosa, "Towards automatic image enhancement using Genetic Algorithms," 2000, pp. 1535-1542.
- [108] S. Hashemi, S. Kiani, N. Noroozi, and M. E. Moghaddam, "An image contrast enhancement method based on genetic algorithm," *Pattern Recognition Letters*, vol. 31, pp. 1816-1824, 2010.
- [109] B. Bhanu, S. Lee, and J. Ming, "Adaptive image segmentation using a genetic algorithm," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 25, pp. 1543-1567, 1995.
- [110] W. B. Tao, J. W. Tian, and J. Liu, "Image segmentation by three-level thresholding based on maximum fuzzy entropy and genetic algorithm," *Pattern Recognition Letters*, vol. 24, pp. 3069-3078, 2003.
- [111] J. M. Rouet, J. J. Jacq, and C. Roux, "Genetic algorithms for a robust 3-D MR-CT registration," *IEEE Transactions on Information Technology in Biomedicine*, vol. 4, pp. 126-136, 2000.
- [112] Q. Wang and X. Li, "Application of improved genetic algorithm in practical medical image registration," *International Journal of Digital Content Technology and its Applications*, vol. 5, pp. 60-67, 2011.
- [113] F. Maes, D. Vandermeulen, and P. Suetens, "Medical image registration using mutual information," *Proceedings of the IEEE*, vol. 91, pp. 1699-1721, 2003.
- [114] C. K. Chow, H. T. Tsui, and T. Lee, "Surface registration using a dynamic genetic algorithm," *Pattern Recognition*, vol. 37, pp. 105-117, 2004.
- [115] R. He and P. A. Narayana, "Global optimization of mutual information: Application to three-dimensional retrospective registration of magnetic resonance images," *Computerized Medical Imaging and Graphics*, vol. 26, pp. 277-292, 2002.
- [116] J. Schwiegerling and J. E. Greivenkamp, "Using corneal height maps and polynomial decomposition to determine corneal aberrations," *Optometry and Vision Science*, vol. 74, pp. 906-916, 1997.
- [117] R. J. Noll, "ZERNIKE POLYNOMIALS AND ATMOSPHERIC TURBULENCE," *J Opt Soc Am*, vol. 66, pp. 207-211, 1976.
- [118] D. R. Iskander, M. J. Collins, and B. Davis, "Optimal modeling of corneal surfaces with Zernike polynomials," *IEEE Transactions on Biomedical Engineering*, vol. 48, pp. 87-95, 2001.
- [119] M. Szkulmowski, I. Gorczynska, D. Szlag, M. Sylwestrzak, A. Kowalczyk, and M. Wojtkowski, "Efficient reduction of speckle noise in Optical Coherence Tomography," *Optics Express*, vol. 20, pp. 1337-1359, 2012/01/16 2012.
- [120] B. R. Davidson and J. K. Barton, "Application of optical coherence tomography to automated contact lens metrology," *Journal of Biomedical Optics*, vol. 15, pp. 016009-016009-9-016009-016009-9, 2010.
- [121] T. Dietenbeck, M. Alessandrini, D. Friboulet, and O. Bernard, "Creaseg: A free software for the evaluation of image segmentation algorithms based on level-set," 2010, pp. 665-668.

- [122] A. Fitzgibbon, M. Pilu, and R. B. Fisher, "Direct least square fitting of ellipses," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, pp. 476-480, 1999.
- [123] I. M. Gelfand, I. M. Gelfand, and S. V. e. Fomin, *Calculus of variations*: Courier Corporation, 2000.
- [124] T. Chan and W. Zhu, "Level set based shape prior segmentation," 2005, pp. 1164-1170.
- [125] D. Cremers, N. Sochen, and C. Schnörr, "A multiphase dynamic labeling model for variational recognition-driven image segmentation," *International Journal of Computer Vision*, vol. 66, pp. 67-81, 2006.
- [126] D. Cremers, N. Sochen, and C. Schnörr, "Towards recognition-based variational segmentation using shape priors and dynamic labeling," vol. 2695, ed, 2003, pp. 388-400.
- [127] S. Radhakrishnan, J. See, S. D. Smith, W. P. Nolan, Z. Ce, D. S. Friedman, *et al.*, "Reproducibility of anterior chamber angle measurements obtained with anterior segment optical coherence tomography," *Investigative Ophthalmology & Visual Science*, vol. 48, pp. 3683-3688, 2007.
- [128] D. Freedman and T. Zhang, "Interactive graph cut based segmentation with shape priors," 2005, pp. 755-762.
- [129] M. Shen, L. Cui, M. Li, D. Zhu, M. R. Wang, and J. Wanga, "Extended scan depth optical coherence tomography for evaluating ocular surface shape," *Journal of Biomedical Optics*, vol. 16, 2011.
- [130] E. Bae and X. C. Tai, "Efficient global minimization for the multiphase Chan-Vese model of image segmentation," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* vol. 5681 LNCS, ed, 2009, pp. 28-41.
- [131] J. A. Eichel, K. K. Bizheva, D. A. Clausi, and P. W. Fieguth, *Automated 3D reconstruction and segmentation from optical coherence tomography* vol. 6313 LNCS, 2010.