Characterization of Retinal OCT images with macular holes

Jakob Thomadsen

Kongens Lyngby 2006
IMM-Master Thesis-2006-11

Technical University of Denmark
Informatics and Mathematical Modelling
Building 321, DK-2800 Kongens Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk

Abstract

An imaging technology called Optical Coherence Tomography (OCT) has among other places found its application within ophthalmology. OCT can produce high-resolution cross sectional images of the internal microstructure of the retina. In this thesis methods to reduce noise in OCT will be investigated. OCT images are in particular affected by a type of noise called speckle that arise due to interference. Three different types of diffusion are applied to single OCT images to test its ability to reduce noise. None of the diffusion methods produce satisfactory results, so an iterative method is developed that averages images taken of the same retinal location. Each image is registered vertically and horizontally to a template, before averaging is done. The method is robust to parametrical changes, and the average image has significantly less noise than the originals.

Retinal OCT images taken of a pathology called macular hole, are investigated to estimate descriptive parameters that could be relevant in evaluating the current state of the pathology. Different descriptors are evaluated pre- and postoperative. These descriptors are to be used in a case study at Herlev Hospital, where different surgical techniques to treat macular hole are evaluated. The descriptors can be extracted once a set of transitional layers have been located. They are found automatically or semi-automatically. If these layers are determined in OCT images scanning the retina at different locations, the neuroretinal thickness can be represented as a surface map or 3D surface, in this way visualizing the entire retina instead of slices of it.

Key words: OCT, optical coherence tomography, speckle, macular hole, regularized dynamic programming, retinal layers, image registration, diffusion, active contours, snakes

Resumé

En visualiseringsteknologi kaldet optisk kohærens tomografi (OCT) har blandt andet fundet sin anvendelse inden for oftalmologien. OCT kan producere højopløselig tværsnitsbilleder af den indre mikrostruktur af nethinden. I denne rapport vil metoder til at reducere støj i OCT blive undersøgt. OCT billeder er i særdeleshed påvirket af en type støj kaldet speckle, som opstår på grund af interferens. Tre forskellige typer diffusion er anvendt på enkelte OCT billeder for at teste deres evner til at reducere støj. Ingen af diffusionsmetoderne giver et tilfredsstillende resultat, så en iterativ metode er udviklet, som producerer et gennemsnit af billeder taget af den samme nethinde position. Hvert billede er registreret vertikalt samt horisontalt til en template, før gennemsnittet er taget. Metoden er robust overfor ændringer i parametrene, og gennemsnitsbilledet har signifikant mindre støj end de oprindelige.

OCT billeder taget af en patologi kaldet maculahul er undersøgt for at estimere deskriptive parametre, som kan være relevante til at evaluere stadiet af patologien. Forskellige deskriptorer er evalueret præ- og postoperativt. Disse deskriptorer skal anvendes i et casestudie på Herlev Amtssygehus, hvor forskellige operationsteknikker til at behandle maculahul undersøges. Deskriptorerne kan bestemmes når et sæt af overgangslag er lokaliseret i billedet. Lagene er fundet automatisk eller semiautomatisk. Hvis disse lag er lokaliseret i OCT billeder, der skanner nethinden forskellige steder, kan nethindens tykkelse repræsenteres som et tykkelseskort eller 3D-overflade, og på den måde visualisere hele nethinden i stedet for tværsnit af den.

Nøgleord: OCT, optisk kohærens tomografi, speckle, maculahul, regulariseret dynamisk programmering, nethindelag, billedregistrering, diffusion, aktive konturer, snakes

Preface

This thesis has been prepared at the section for Image Analysis, Department of Informatics and Mathematical Modelling, at the Technical University of Denmark in partial fulfillment of the requirements for acquiring the degree Master of Science in Engineering. The work has been done in participation with the Optics and Plasma Research Department at Risø National Laboratory and the Department of Ophthalmology at Herlev Hospital.

Currently a paper, based on the methods to reduce noise presented in this thesis, is under preparation to be submitted to Optics Express.

Lyngby, February 2006

Jakob Thomadsen

Acknowledgements

Even though this thesis is written as an individual project, it would never have been the same without the influence from the following people.

From DTU I would like to thank my supervisor Bjarne K. Ersbøll and co-supervisor Rasmus Larsen first of all for introducing me to image analysis through my bachelor project, but also for their encouragement and support throughout the course of the project. All the people at the office is thanked for providing a fruitful work environment, where people took the time to help each other.

My supervisor from Risø Thomas M. Jørgensen is thanked for letting me bring to reality and extend some of his ideas to improve the quality of OCT. The hours spend discussing different aspects of the project has definitely made the thesis much better. I would also like to thank Peter Andersen from Risø for introducing me to the world of OCT and its applications.

For an inspiring collaboration and explaining me the medical aspects I would like to thank Ulrik Christensen and Kristian Krøyer from Herlev Hospital, and likewise Birgit Sander, who I also thank for helping me write scholarship applications.

JC Jacobsen scholarship and Bernhard Eisenreich Sandersen and wife scholarship is thanked for the financial support, so I did not have to work and could focus solely on my project.

And finally I thank my girlfriend Luise Gade for all the love and support you have given me.

Contents

Chapter 1

Introduction

For a number of serious eye diseases, the pathological changes are localized in the retina, ie. the back of the eye where the light is focused. This is the case for diabetes and glaucoma, but also less known diseases such as macular hole, where a rupture happens in the central part of the retina. Since this is where our sharpest vision is localized, a macular hole leads to a significant loss of vision, including the ability to read on the affected eye. It is estimated that the disease affects 1500 Danes a year [1].

At the ophthalmological department at Herlev Hospital there is a center for surgical treatment of macular hole. It is one of the leading departments in Denmark with approximately 100 surgeries a year.

Surgery is currently the only way to treat a macular hole that does not close spontaneously. It was discovered in 1999 as reported in [2] that peeling of the internal limiting membrane (ILM) on the retina has been found to be a way to stimulate the wound healing. It is an extremely difficult procedure, but it is made easier for the surgeons if the membrane is stained with the dye indocyanine green (ICG). There are concerns that ICG may be toxic to the retina, but on the other hand if the ILM is not peeled, there is a significantly higher risk of the hole not closing. [3].

The rate of hole closures is about 80%, when no peeling is done, and close to 100% when the ILM is peeled. Many experience an improvement in the visual acuity, but for a part this does not happen, even though the hole has been closed. The reason for this discrepancy between the anatomical and functional result has not been understood and there is a need for predicting who will benefit from an operation.

It is imperative that the surgeons have all relevant information available about the patients condition before deciding whether or not to operate. This is where a method called OCT finds its application. OCT stands for "Optical Coherence Tomography" and is a technique that was developed in the early nineties, first reported in [4]. It is an imaging technology that produces high-resolution cross sectional images of the internal microstructure of living tissue. OCT has many similarities with ultrasonic imaging, with a major difference in source, where coherent light is used instead of ultrasound.

The OCT system currently used at Herlev Hospital is a third generation product from Zeiss, called StratusOCT. All OCT images in this thesis have been taken with this system.

Two Ph.D. students at Herlev Hospital are currently performing a randomized case study of different surgical techniques when treating a macular hole. Half of the patients will have the ILM peeled, and the other half will not. A significant evaluation method is OCT. To quantitatively be able to evaluate the changes occurring in the retina, it is relevant to look at several measurable descriptors in the OCT images. Deciding which descriptors could be interesting has been done in collaboration with the two Ph.D. students at Herlev Hospital. As a part of this thesis, software has been developed that finds these relevant descriptors automatically or semi-automatically. This software is to be used as a tool for these Ph.D. students during their following data analysis.

When making decisions based on an OCT image, it is of course beneficial to have as little noise as possible in the image. This can reduce the uncertainty when classifying a given pathology, but also lead to the visualization of minute details and new insight. The amount of noise can be reduced even without altering the system. This has been seen in a method developed at Risø based on averaging several OCT images [5]. This way a type of noise known as speckle is significantly reduced. This method does not work satisfactory on OCT images of macular holes, so a new method to reduce speckle will be developed. Two approaches will be taken to achieve this. One being processing a single image, and thereby improving the quality, the other being an extension of the method developed at Risø, where several images of the same area are combined to produce a single image.

The primary goal for the two Ph.D students at Herlev Hospital is to evaluate the surgical results achieved with ICG assisted ILM peeling, vs. no peeling. But during the case study insight about the pathogenesis of macular hole may be gained. This could for instance be about which macular holes spontaneously regress and which does not. Then surgery could be performed earlier in the ones not expected to regress, and later in the ones that may regress.

1.1 Thesis Overview

This thesis is separated into four natural parts.

Part I - Background introduces OCT, the eye and dynamic programming

Part II - Image Enhancement describes the enhancement of OCT images by use of diffusion and averaging of multiple images

Part III - Applications describes the extraction of the pre- and postoperative descriptors and the interpolation of the retinal thickness

Part IV - Discussion sums up the thesis

1.2 Nomenclature

The most significant variables used in this thesis are listed below:

$d(i, j)$	Distance to (i, j), ie. minimum length of paths ending at (i, j)
E_{ext}	External energy of snake
E_{int}	Internal energy of snake
E_{snake}	Energy of snake, ie. sum of E_{ext} and E_{int}
$I(i, j, t)$	Image - row, column, time
M	Number of columns
N	Number of rows
p	Order of path, ie. max. vertical change between adjacent points
P	Path through an image
$R_{a_i b_j}(m)$	Raw cross-correlation of column i in a and j in b with shift m
Vol_i	Estimate of volume of macular hole based on i OCT scans
λ	Regularization parameter in regularized dynamic programming

1.3 Abbreviations

All abbreviations used in this thesis are listed below:

CNR	Contrast to Noise Ratio
ICG	Indocyanine Green
ILM	Internal Limiting Membrane
IS/OS	Inner and Outer Photoreceptor Segments
OCT	Optical Coherence Tomography
Pre-OMaH Tool	Pre-Operative Macular Hole Tool
Post-OMaH Tool	Post-Operative Macular Hole Tool
RNFL	Retinal Nerve Fiber Layer
RPE	Retinal Pigment Epithelium
SNR	Signal to Noise Ratio

Part I

Background

Chapter 2

Optical Coherence Tomography

The development of the laser in the 1960's gave the physicians a new surgical instrument. But most of the applications so far have not taken advantage of the coherent properties of the laser, ie. it emits photons with the same wavelength, phase and direction. Instead the laser is often used as a lighting or heating source [6]. There are exceptions to this, and one of these is optical coherence tomography (OCT), a high resolution cross-sectional imaging method. It has the abilities to achieve a probing depth exceeding 2cm in transparent tissue, such as the eye, and it is possible to visualize structure 1-2mm beneath the surface in highly scattering tissues like skin [7]. This is possible with a transverse and longitudinal resolution of a few micrometers [4].

The technique was extended for the first time from the 1D case of optical coherence-domain reflectometry to 2D in 1991 as reported in [4]. Lately it has been extended to the 3D case [8],[9], but the 3D-method has not been considered ready for market, since no products have been launched yet.

A schematic OCT scanner is shown in figure 2.1. The central part is a fiber optic Michelson interferometer, which is illuminated by a low-coherent, broadband light source. The sample is placed at the end of one interferometer arm, and a reference mirror at the other. The source field is split into a sample field E_s

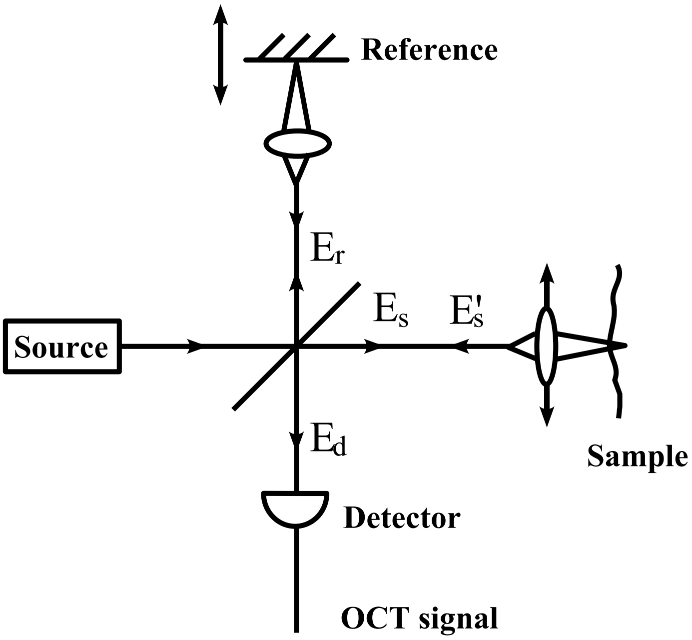

Figure 2.1: Schematic OCT scanner. The source light is split into a sample and a reference beam. Reflections from the two are combined and detected by a photodiode.

and a reference field E_r. After scattering back from some point in the sample, the modified sample field E_s' mixes with E_r in the detector. This intensity J_d depends linearly on the real part of the cross-correlation between E_s' and E_r.

$$J_d = \langle |E_d(\tau)|^2 \rangle = 0.5(J_s' + J_r) + Re\langle E_r^*(t+\tau)E_s'(t)\rangle \tag{2.1}$$

Where J_s' and J_r are the mean intensities returning from the sample and reference arm of the interferometer. The first term on the right hand side in equation 2.1 is of no interest, but the second term depends on the optical time delay τ set by the position of the reference mirror. It is the term that carry information about the tissue structure. The interference patterns formed depends on the temporal and spatial coherence of E_s' and E_r. The interferometer functions as a cross-correlator, and the amplitude of the interference signal after integration on the detector provides a measure of the amplitude of the cross-correlation, ie. it functions as a measure of a high reflection or back-scattering at the current depth.

When the light source satisfies the quasimonochromatic condition, ie. its center frequency greatly exceeds its bandwidth, the correlation amplitude can be expressed as

$$Re\langle E_r^*(t+\tau)E_s'(t)\rangle = |G(\tau)|cos(2\pi\nu_0\tau + \phi(\tau)) \tag{2.2}$$

Where $|G(\tau)|$ and $\phi(\tau)$ are respectively the argument and phase of the complex temporal coherence and ν_0 is the center frequency of the source. The complex temporal coherence does besides the sample also depend on the spectral shape of the source. Broadband sources are desirable in OCT systems, since they produce interference patterns of short temporal and spatial extent[6].

The OCT scanner performs multiple longitudinal scans at a series of lateral locations to provide a 2D map of reflection sites in the sample. This is analogous to ultrasonic pulse-echo imaging (B-mode scanning), with a major difference in choice of source. This analogy has caused an analogy in terminology. An OCT image is called a B-scan, and a single longitudinal scan is referred to as an A-scan.

2.1 Speckle

As previously mentioned OCT is a method that relies on measuring the temporal and spatial coherence of the signal reflected or backscattered from the sample. In a highly scattering sample such as tissue, the coherence is also what gives rise to speckle, a type of noise that reduce contrast and make boundaries between highly scattering structures difficult to distinguish. The noise arise from interference between coherent waves backscattered from nearby scatterers in the sample.

An interference pattern showing speckle is shown in 2.2(a). The image shows the reflection from a stationary titanium block. An image with limited speckle is shown in 2.2(b). The image is taken of the same block, but this time the block was in motion. When the block is moving, the speckle pattern is changing, and the 0.1s integration time has reduced the effect significantly.

The effect of speckle is well known in ultrasonic and radar imaging. It occurs when waves from a coherent source encounters scatters separated by distances near that of the coherence length of the source. Speckle is not truly a noise in the typical sense. It is signal dependent and it carries information about the sample being imaged.

Explained in greater detail, the intensity detected is a function of both the phase and the amplitude of the complex temporal coherence, as was seen in equation 2.2. It is the sensitivity to the phase that makes OCT susceptible to

(a) Significant speckle

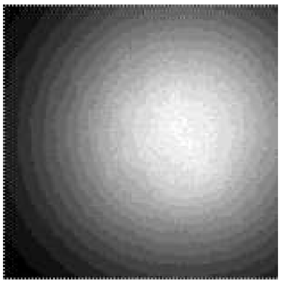

(b) Limited speckle

Figure 2.2: (a) shows a speckle filled interference pattern. The image shows the reflection from a stationary titanium block. (b) shows very limited speckle. The image is taken of the same block, but this time the block was in motion. When the block is moving, the speckle pattern is changing, and the 0.1s integration time has reduced the effect. From [10].

the effect of speckle. When the sample is a tissue containing several layers of scatterers $\phi(\tau)$ can no longer be treated as a deterministic variable because the spatial coherence is lost through random scattering [11].

The shape of a coherent wave entering a tissue and being scattered back, will be affected by multiple forward scattering on the way forward and back and multiple backscattering in the tissue at the probing distance. Both of these scattering types effects the shape of the returning wave, and create localized regions of constructive and destructive interference, seen as speckle in the image.

These two effects are what have been believed to be the primary reason for speckle formation in OCT images. A study [12] indicates that multiple scattered light loses its coherence, such that the speckle formation mostly comes from the superposition of several slightly scattered waves. No matter which part is the dominating, the effect is obvious.

Ways to reduce the effects of speckle can be split into polarization diversity, frequency compounding, spatial compounding and image postprocessing. The first three methods are not possible without changing the apparatus. Some of the image postprocessing methods that have been applied are median filter, iterative deconvolution methods [13], wavelet filtering [14], [15], rotating kernel transformation [16] and other adaptive smoothing methods [17].

Another way of suppressing the effect of noise is to average a number of measurements, the underlying assumption being that the noise is stochastically varying in the measurements taken. This would be equivalent to spatial compounding, without having control of the minute shifts occurring between each scan. However as was seen in figure 2.2, a physical change of the apparatus relative to the sample is needed to change the speckle pattern. This is why taking two images consecutively and averaging these will reduce the effect of speckle more than taking two A-scans at a time and averaging these, since minimal movement of the sample occurs in the brief time between two A-scans are taken, and no movement of the apparatus has happened.

Chapter 3

The Eye

The human eye is the organ that gives us the sense of sight. Most of the structure is devoted to the one task of focusing light onto the retina. All of the individual components through which light travels within the eye before reaching the retina are transparent, minimizing dimming of the light. The cornea and lens focus light rays onto the retina. This light causes chemical changes in the photosensitive cells of the retina, the products of which trigger nerve impulses that travel to the brain.

The following section has been written mostly based on a text from [18]. When reading this section, it can be beneficial to look at figure 3.1 to visualize where the different parts of the eye are located. Light enters the eye from an external medium, passes through the cornea, into the anterior chamber filled with water. Most of the light refraction occurs at the cornea which has a fixed curvature. The iris, between the cornea and the lens, is a colored ring of muscle fibres. Light must pass though the center of the iris, the pupil. The lens is a convex shaped disk which focuses light onto the retina. The cavity of the eye behind the lens is called the vitreous body. It is a clear gelatinous substance that lets light through to the retina.

Wrapped around the vitreous body are three layers of tissue that maintains the shape of the eye. The outermost is the sclera which gives most of the eye its white color. It protects the inner components of the eye. On the inner side of

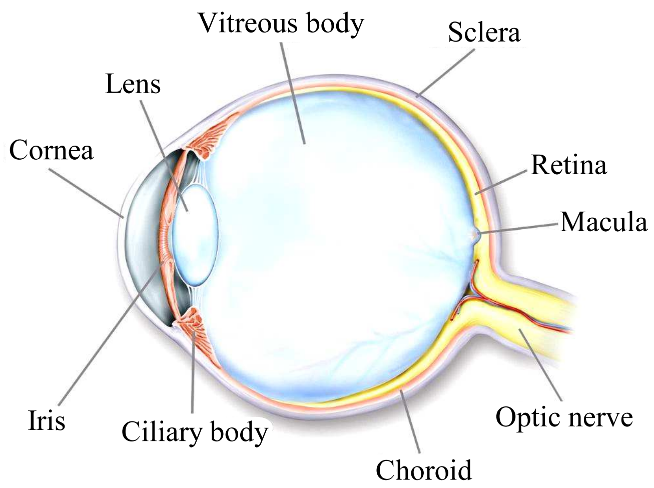

Figure 3.1: Cross-sectional view of the human eye. Adapted from a drawing from Øjenforeningen.

the sclera is the choroid which contains blood vessels that supply the retinal cells with necessary oxygen and removes the waste products of respiration. The choroid gives the inner eye a dark color, which prevents disruptive reflections within the eye. The innermost layer of the eye is the retina, containing the photosensitive rod and cone cells, and neurons. The retina is less than 0.5mm thick.

The retina is a relatively smooth layer. It does however have two points at which it is different: the optic nerve head and the fovea. The optic nerve head is the point on the retina where the optic nerve pierces the three previously named layers to connect to the nerve cells on the top of the retina. No photosensitive cells exist at this point. The fovea is the center of the macula lying directly opposite the lens. The fovea is a dip in the retina and is densely packed with cone cells. It is largely responsible for color vision in humans, and enables high acuity that eg. is necessary when reading.

The retina can be visualized with OCT. An example can be seen on figure 3.2. The width of the image is 6mm and the height is 2mm. It is a straight line scan through the fovea, the cup in the middle of the image. It is exported from StratusOCT, and no postprocessing has been done except for the colormap, ranging from dark blue, where no light is reflected to red with high reflectance. A few of the layers that can be seen in the image have been marked with arrows. Starting from the top, the retinal nerve fiber layer (RNFL) can be seen

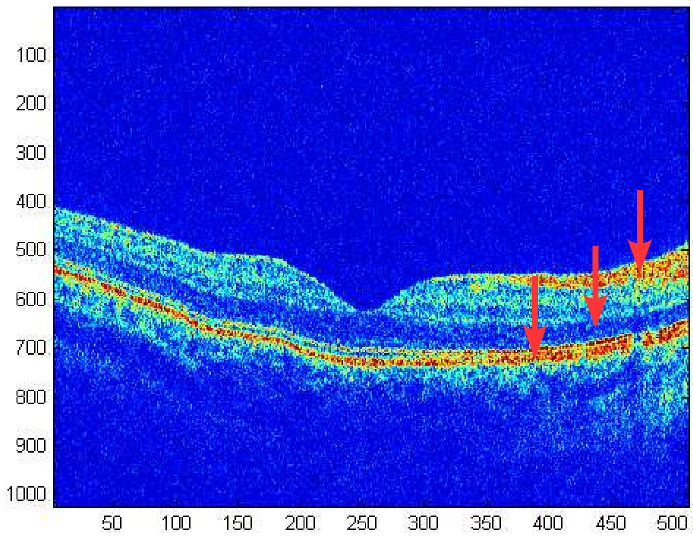

Figure 3.2: Straight line scan through the fovea of an individual with no pathologies in the eye. Arrows indicate, starting from the top: RNFL, outer nuclear layer and RPE.

distinctly on the right side of the image. This is because the optic nerve head is located on this side, and all the nerve fibers runs to the optic nerve head. The hyporeflective layer marked with the second arrow is the outer nuclear layer consisting of the cell bodies of rods and cones. The hyperreflective layer marked with the last arrow is actually two layers, the bottom and most prominent one being the retinal pigment epithelium (RPE), marking where the neuroretina ends, and the top one is believed to be the junction between the inner and outer photoreceptor segment (IS/OS). See eg. [19] for a further discussion of visualizations in retinal OCT imaging.

3.1 Macular Hole

For a number of serious eye pathologies, the changes are localized in the retina. As examples can be mentioned diabetes and glaucoma, but also less known pathologies such as macular hole. It is a disease where a rupture happens in the macula, and a hole down to the RPE is formed. An OCT image of a macular hole can be seen in Figure 3.3.

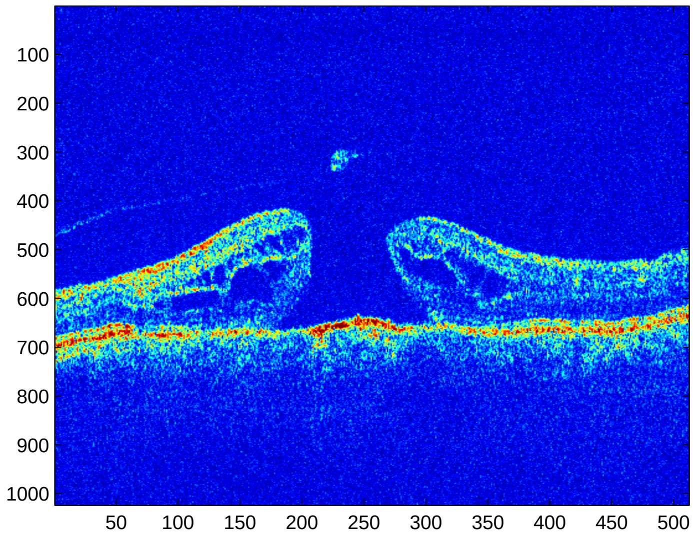

Figure 3.3: Straight line scan through the fovea of an individual with a macular hole. The photoreceptor layer has ruptured in the fovea, and separated from the RPE.

Since the macula is where our sharpest vision is located, a macular hole gives a significant loss of vision, including the ability to read on the affected eye. It is estimated that the disease affects approximately 1500 Danes a year [1]. In rare cases a macular hole can develop because of trauma, but in most cases they are idiopathic. Since the early nineties it has been possible to treat a macular hole with advanced surgery, but a part of the patients does not regain their visual acuity.

A macular hole develop through several stages, starting with an impending hole. About half of the impending holes regress spontaneously, and the other half progress to full thickness macular holes if not treated. The macular hole stages were originally hypothesized by Gass [20] in 1988. Since the development of OCT, Gass' theory has been revised. The currently believed development of a macular hole is shown in figure 3.4. The diagrams are copied from [1]. A sketch of a normal fovea as seen from an OCT is shown in diagram A. The border of the vitreous body is drawn as a dotted line. An impending hole can be seen in diagram B. There is perifoveal detachment of the vitreous body, and a cystic space in the inner part of the fovea. Diagram C shows a stage 2 macular hole. A fully developed stage 3 macular hole can be seen in diagram D. The roof of the foveal cyst is suspended in front of the macular hole. Diagram E shows a stage 4 macular hole with complete posterior vitreous detachment. In the early stages the individual does not notice the pathology, but if the condition gets worse, the individual will experience image distortion and the ability to read will be affected.

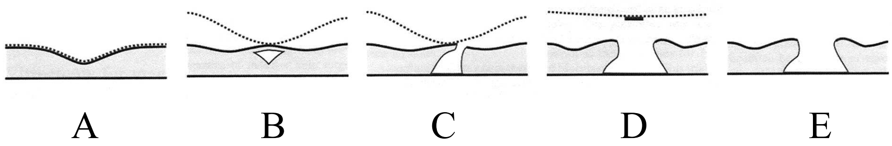

Figure 3.4: Typical macular hole development. Diagram A: Normal fovea as seen from an OCT. The border of the vitreous body is drawn as a dotted line. Diagram B: An impending hole can be seen. There is perifoveal detachment of the vitreous body, and a cystic space in the inner part of the fovea. Diagram C: Stage 2 macular hole. Diagram D: A fully developed stage 3 macular hole. The roof of the foveal cyst is suspended in front of the macular hole. Diagram E: Stage 4 macular hole with complete posterior vitreous detachment. Copied from [1].

The etiology of macular holes is not completely known, but it is believed that it is due to traction in the vitreous body on the fovea and some kind of degenerative dissolution of the inner retinal layers in the fovea [21]. Recently a hydrodynamic model that balances fluid flow in the hole with fluid pumped across the RPE has been suggested [22]. It implies that the relief of traction alone is not sufficient in closing macular holes, unless it is accompanied by a reduction in hole diameter. This is currently achieved by injection of a gas in the back of the vitreous that applies pressure on the macula.

The loss of visual acuity when having a macular hole, is because the photoreceptors in the fovea are missing. There are indications that the missing photoreceptors are not lost, but have been moved peripherally away from the hole, where they still function, but with reduced effect. This causes the photoreceptors to project the visual field abnormally, so the affected individual will experience a severe distortion as well as a reduced visual acuity of everything seen at the central field of vision.

3.1.1 Treatment

An operation is currently the only way to treat a macular hole that does not close spontaneously. The aim is to close the hole and hopefully thereby improve the visual acuity and reduce the distortion. The operation requires removal of the vitreous lying in front of the macular hole. It was discovered in 1999 as reported in [2] that peeling of the ILM on the retina has been found to be a way to stimulate the wound healing in the macula. It is an extremely difficult procedure, but it is made easier for the surgeons if the membrane is stained with the dye ICG. Currently, ICG assisted ILM peeling is the preferred method for stage 3 and 4 macular holes in many Scandinavian vitreoretinal centres [1]. But there are concerns that ICG dye may be toxic to the RPE, but on the other hand if the ILM is not peeled, there is a significantly higher risk of the hole not closing. [3].

At the end of surgery, the back of the vitreous is filled with an inactive gas that is absorbed over the next couple of weeks. Currently the operation must be followed by a period of several days, where the individuals face must be kept facing downwards several hours a day, so the gas filling presses against the macula.

The rate of hole closures is about 80%, when no peeling is done and close to 100% when the ILM is peeled. Many experience an improvement in the visual acuity, but for some this does not happen, even though the hole has closed. This discrepancy can not be explained and there is a need for predicting who will benefit from an operation and when it should be performed.

The treatment of macular holes has developed rapidly since surgeons have started treating it less than 20 years ago, but since the pathology is not completely understood, there might be basis for improvements of the treatment over the coming years. These improvements will come from insight that may be gained from studies such as the one being performed at Herlev Hospital.

Chapter 4

Dynamic Programming

Dynamic programming is a method for finding the shortest path between two sets of pixels. Its general principles was first introduced by Bellman in [23]. Some of the advantages of dynamic programming are that it is simple to program, efficient computationally and finds the optimal path. But it has the shortcoming of not allowing any smoothness constraints on its path. An extension was introduced in [24], which applied a penalty depending on the shape of the contour. To maintain optimality it was an iterative process. The method was adapted by Buckley in [25], where a method for finding smooth shortest paths across images was introduced. In the paper the method is referred to as "regularised shortest-path extraction", but I will refer to it as "regularized dynamic programming".

First the standard dynamic programming algorithm will be described. Dynamic programming has been extensively used for finding the shortest path on a rectangular grid. The shortest path is defined as the path that minimizes the sum of the weights of the arcs on the path. The path must always go from one layer to the following layer. Explained on an image, it means that the path must move from one column to the next in the case where we are interested in a solution from one side to the other. The procedure is symmetric in the sense that a left-to-right search gives the same result as a right-to-left, when the shortest path is unique.

For every pixel in an image with N rows and M columns, let $I(i,j)$ be the grey value of the pixel located at (i,j). A path P of order p from the left side of the image to a point on the right side of the image (L,M) is a set of M pixels

$$\begin{aligned} &P = \{(i_1,1),\ldots,(i_M,M)\}, \\ &\text{with } 1 \le i_c \le N \text{ for } c = 1,\ldots,M, \\ &\text{and } |i_{c+1} - i_c| \le p, \\ &\text{for } c = 1,\ldots,M-1. \end{aligned} \tag{4.1}$$

The last pixel must of course be the point, so $i_M = L$. Explained in words, P is a set of pixels, one pixel in each of the M columns, with the last one being (L,M). The points have pth-order connectivity. If $p = 1$, the path is an 8-connected set. The length of a path P is defined as the sum of the pixelvalues in the path

$$\sum_{c=1}^{M} I(i_c,c). \tag{4.2}$$

The shortest path going from one side of the image to the other is the path that minimizes equation 4.2. To help find this path a distance $d(i,j)$ is introduced for every (i,j) as the minimum length of all paths going from the left side of the image and ending at (i,j). For the trivial case $j = 1$ there is only one path ending at the pixel, the single pixel path $P = \{(i,1)\}$, so

$$d(i,1) = I(i,1), \quad \text{for all } i = 1,\ldots,N. \tag{4.3}$$

This is extended to other columns recursively by use of the connectivity constraint. For a given position (i,j) the path must come from one of the $2p+1$ closest points located on the previously column, ie. $(i+k,j-1)$ where $-p \le k \le p$. This is true for all $j \ge 2$. The minimum distance to reach a given point can therefore be written as

$$d(i,j) = I(i,j) + \min_{k:|k|\le p} d(i+k,j-1), \tag{4.4}$$

which along with the initial condition, equation 4.3, enables a recursive calculation of $d(i,j)$ for all i and j. The value of k that leads to a minimum should be stored for later backtracking of the optimal path

$$k(i,j) = \operatorname*{argmin}_{k:|k|\le p} d(i+k,j-1). \tag{4.5}$$

Once the right side of the image is reached, the length of the shortest path can be found

$$\min_{1\le i\le N} d(i,M), \tag{4.6}$$

and the last point on the path is the point $(\hat{i}_M,M)$ that achieves the minimum. The path that led to this minimum value can now be backtracked starting with

$\hat{i}_{M-1} = \hat{i}_M + k(\hat{i}_M, M)$. This can be repeated for every column $\hat{i}_{j-1} = \hat{i}_j + k(\hat{i}_j, j)$ until the left side of the image is reached and the entire path is known. The computational cost of the algorithm is $O((2p+1)NM)$.

The standard method is tested on a section of a concentric shaped OCT image taken around the fovea. We are not interested in a particular path, just a continuous path that follows the dark parts of the image the best. Figure 4.1(a) shows the original image, and figure 4.1(b) shows the path found with dynamic programming with $p = 1$. For higher values of p, the path becomes more rough, since the continuity constraint is reduced. The path is already fairly rough and in this implementation there is no way of constraining with respect to smoothness. If the algorithm is extended to regularized dynamic programming this can be achieved.

4.1 Regularized Dynamic Programming

In regularized dynamic programming a penalty that is proportional to a measure of the roughness of the path is added to the length of the path. The roughness of a path $P = \{(i_1, 1), \ldots, (i_M, M)\}$ is defined as [25]

$$\sum_{c=2}^{M-1} (i_{c-1} - 2i_c + i_{c+1})^2. \tag{4.7}$$

This leads to a new definition of the length of path P

$$\sum_{c=1}^{M} I(i_c, c) + \lambda \sum_{c=2}^{M-1} (i_{c-1} - 2i_c + i_{c+1})^2, \tag{4.8}$$

where λ is a regularization constant. The optimal path is still where the length is minimized.

If $d(i,j)$ was defined the same way as for standard dynamic programming, optimality would not be guaranteed since it is not certain that the path, which minimizes the length to a given point will minimize a longer path, since its total roughness contribution is unknown. Therefore we need to define $d(i,j,k)$ as the length of the shortest path to reach point (i,j) via $(i+k, j-1)$, for $k = -p, \ldots, p$.

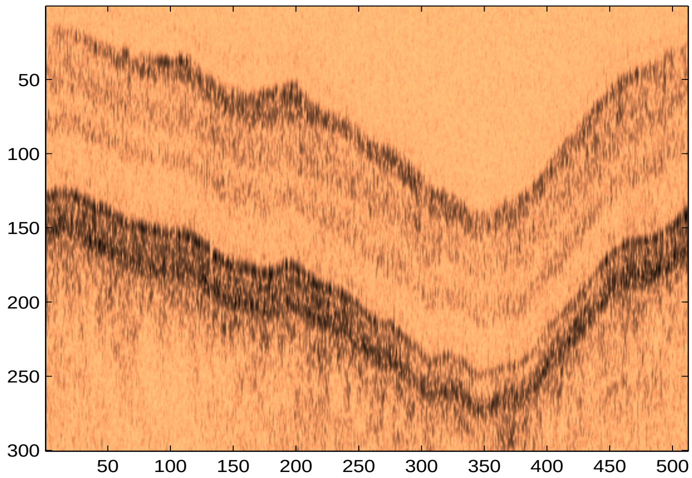

(a) Original image

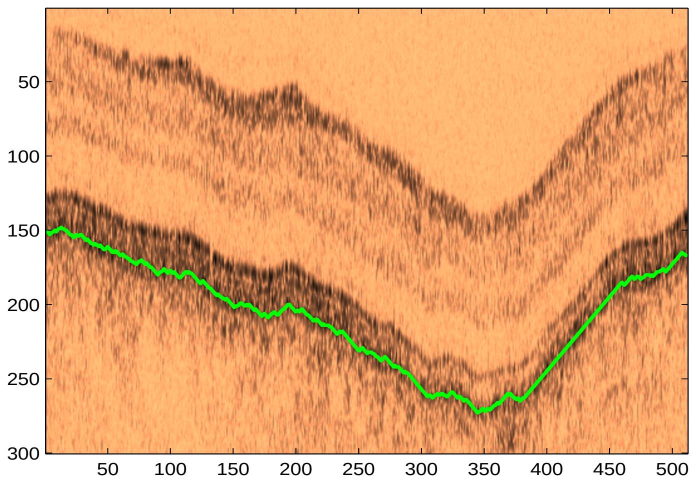

(b) Dynamic programming with $p = 1$

Figure 4.1: A concentric scan around the fovea is used for testing dynamic programming vs. regularized dynamic programming. The original image is shown in (a). The shortest path found with dynamic programming with $p = 1$ is shown in (b). The path is fairly rough and there is no way of constraining with respect to smoothness.

The new initialization equation is

$$d(i, 2, k) = I(i, 2) + I(i + k, 1), \tag{4.9}$$

which holds for $i = 1, \ldots, N$ and $k = -p, \ldots, p$ when the pixel positions are within the allowed intervals.

A path from (i, j) via $(i + k, j - 1)$ must pass through $(i + k + l, j - 2)$ for some $-p \leq l \leq p$. This gives a roughness contribution of $\lambda(l - k)^2$. The extension of the recursion formula is then going to be

$$d(i, j, k) = I(i, j) + \min_{l:|l| \leq p} (d(i + k, j - 1, l) + \lambda(l - k)^2), \tag{4.10}$$

with which it is possible to calculate all values of $d(i, j, k)$. Again the values that leads to a minimum is stored for later backtracking

$$l(i, j, k) = \underset{l:|l| \leq p}{\operatorname{argmin}}(d(i + k, j - 1, l) + \lambda(l - k)^2). \tag{4.11}$$

When backtracking, the last point and the direction it came from is found as

$$(\hat{i}_M, \hat{k}_M) = \underset{(i,k)}{\operatorname{argmin}}\, d(i, M, k). \tag{4.12}$$

The last point on the optimal path is, $(\hat{i}_M, M)$. The second to last point is also known $(\hat{i}_{M-1}, M - 1)$, where $\hat{i}_{M-1} = \hat{i}_M + \hat{k}_M$. The rest of the points on the optimal path can be calculated as $\hat{i}_{j-1} = \hat{i}_j + \hat{k}_j$ with $\hat{k}_j = l(\hat{i}_{j+1}, j + 1, \hat{k}_{j+1})$. The computational cost of the regularized version of dynamic programming is $O((2p + 1)^2 NM)$.

The regularized dynamic programming algorithm has been tested on the image shown in figure 4.1(a). The result with a fairly large λ, one half of the interval between the maximum and minimum pixelvalues, and $p = 1$ is shown in figure 4.2(a). The apparent roughness of the shortest path is small as expected, since the length is minimized according to 4.8. This does not mean that a smooth path has been found though. A piecewise linear path has been found. This is because the only allowed slopes are 0 and ± 1. A gradual change in slope would be a more desirable result.

Increasing p is not going to improve the result, since still no slopes between 0 and 1 will be allowed. If on the other hand the image is subsampled with a factor q vertically, and p is increased by the same factor, suddenly $q - 1$ slopes between 0 and 1 will be permissible. The maximum allowed slope will stay the same, but the allowed slopes will be subsampled by a factor q. This trick is fairly expensive computationally though. If $p = 1$, the computational cost,

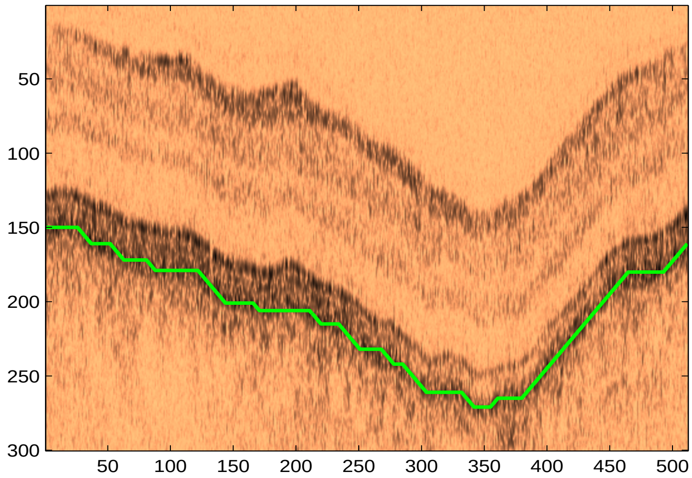

(a) Regularized dynamic programming

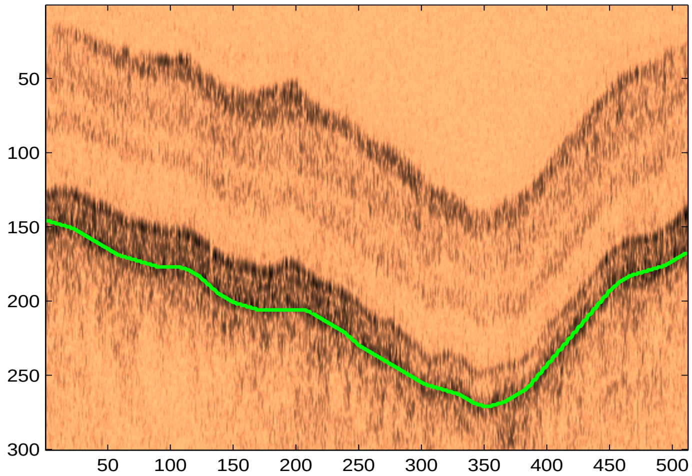

(b) Regularized dynamic programming on changed resolution

Figure 4.2: In (a) regularized dynamic programming with λ equal to half of the interval between the maximum and minimum pixelvalue has been tested on the image shown in 4.1(a). The path is not smooth though, it is piecewise linear. This is because the only allowed slopes are 0 and ± 1. In figure (b) the horizontal resolution has been decreased by a factor 4, before the optimal path was found. The effective λ is the same as in (a). The improvement of the smoothness of the shortest path is significant.

of subsampling vertically by a factor of 2, will increase p and M by a factor of 2, which leads to a computational increase compared with non-subsampled regularized dynamic programming by a factor of $\frac{50}{9} \approx 5.6$. Subsampling with a factor of 4, will increase the computational cost with a factor of 36.

Instead of subsampling vertically with a factor q, a reduction in the horizontal resolution by a factor q, while increasing p by the same factor, will lead to the same increase in allowed slopes. This requires a fairly large horizontal resolution to begin with, compared with the structure you want to find. But if this is the case the method improves the result, with minimum increase in computation time. A reduction in horizontal resolution by a factor 2 will increase p, but also reduce M. If $p = 1$ the increase in computation will for $q = 2$ be a factor $\frac{25}{18} \approx 1.4$.

The result of decreasing the horizontal resolution by a factor 4, is shown in figure 4.2(b). This way four positive slopes less than or equal to one are allowed. The effective λ is maintained by reducing the used λ by a factor of 4. The increase in computation is a factor 2.3 compared with the result in figure 4.2(a), but the improvement of the smoothness of the shortest path is also significant.

Part II

Image Enhancement

CHAPTER 5

Single Image

5.1 A-scan Alignment

During the second or so it takes to record an OCT image, vertical shifts in the image occurs. This is due to minor movements of the individuals eye on the scale of tenth of millimeters. These movements are, as to some degree can be seen on figures 3.2 and 3.3, slowly varying from A-scan to A-scan. The reason that no discontinuities are seen in the images is not that it never happen, but simply that it is common practise for the ophthalmologist to discard the image, and takes a new one instead if it occurs.

With the term "alignment", I will refer to the process of adjusting every A-scan in an image such that it seems as if the structure in the image is horizontal. This is not a true representation of the retina, since it is located on a sphere, but given that the scans are normally 6mm wide, compared with the eye diameter of about 25mm, it is a fair approximation.

It should be noted that the aim of alignment is not to align the top layer. If this was the goal, an active contour could be used to find the top layer as reported in [26] for OCT images. The aim is to reduce the vertical movements that has happened, and align the dominating structure, thereby hopefully making the image look closer to reality.

There are several other reasons for aligning the images before further processing. If we want to apply a filter with horizontal components we need to know the horizontal neighbors or a likely candidate. Another reason is that if we want to measure a relative distance between two points a distance on an aligned image is certainly more accurate than measuring on the original image. A third is that the horizontal structure now lines up, and the problem of registering two images horizontally and vertically can be done independently. Since it is assumed that a horizontal structure exists the vertical registration can be done without having registered the images horizontally. Once this have been achieved, the structure in the images that is not horizontal layers stabilizes the horizontal registration.

There exist several methods for alignment of noisy 1D signals. In [27] and [28] methods based on centroid measurements, normalized integrals and correlation are presented. In [27] the method that gave the best result for all noise levels was complete cross-correlation and for low noise levels the result of an iterative cross-correlation method was very close. If the task at hand is to align two 1D signals the complete cross-correlation method dictates estimating the best possible relative shift for every two signals where one or both of the two is involved, not just for the two at hand. Based on these relative shifts we estimate the relative shift between the two signals that minimize the least squares.

If we want to estimate the relative shift for two adjacent A-scans, and not include any other A-scans, the complete cross-correlation method is based on only one possible relative shift and therefore is the same as just doing cross-correlation between the two A-scans. This is the method adopted for aligning two adjacent A-scans.

I will extend the notation used in chapter 4 for 2D signals such that element (i, j) of image k is expressed as $I_k(i, j)$. In this way, the raw cross-correlation for A-scan or column i of image a and column j of image b is defined as

$$R_{a_i b_j}(m) = \sum_{n=1}^{N-m} I_a(n+m, i) I_b(n, j) \qquad m \geq 0, \tag{5.1}$$

and for negative m

$$R_{a_i b_j}(-m) = R_{b_j a_i}(m) \qquad m < 0. \tag{5.2}$$

Where the relative shift m lies in the interval $-N < m < N$ where N is the number of rows.

To align an image, the cross-correlation between every A-scan and the previous A-scan is maximized. Since the relative shifts are assumed continuous the maximum allowed vertical shift is set to 10 pixels. The procedure for finding the optimal shift m_{opt} is for a given A-scan j in image k therefore

$$m_{opt} = \underset{m}{\operatorname{argmax}}\, R_{k_{j-1}k_j}(m) \qquad \text{for } -10 \leq m \leq 10. \tag{5.3}$$

This is done for every A-scan except the first, $1 < j \leq M$, where M is the number of A-scans in an image. The procedure gave reasonable results, the overall structure was aligned. But the procedure was susceptible to noise, in such a way that A-scans would fluctuate a few pixels up and down. Therefore A-scan $j-1$ was replace by the mean of the previous 10 aligned A-scans, to minimize the small scale variation. The results of aligning the OCT images in figures 3.2 and 3.3 can be seen on figures 5.1 and 5.2.

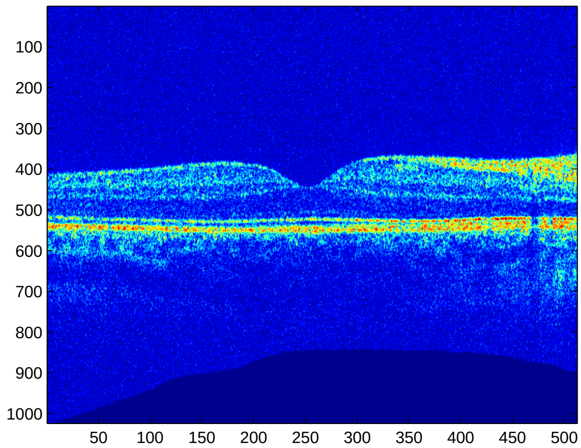

Figure 5.1: The image in figure 3.2 has been aligned by maximizing the cross-correlation between every A-scan and the mean of the aligned previous 10 A-scans. The dominant structure in the image, the RPE, is now horizontal in the image.

It could be argued that when the shift is based on several previous columns, complete cross-correlation can be implemented. But in this case we actually are more interested in aligning the current A-scan to the overall previous structure, not just to the previous noisy A-scan. So in this case this method gives the result we are interested in.

Another way to make the procedure less susceptible to noise could be to apply a low pass filter on every A-scan, but it has not been necessary, since the result is satisfactory.

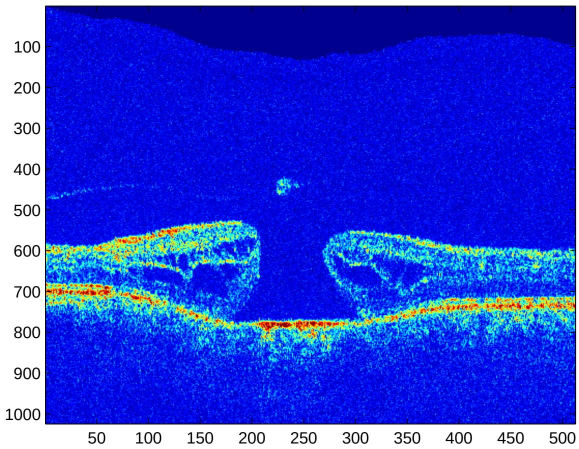

Figure 5.2: The image in figure 3.3 has been aligned by maximizing the cross-correlation between every A-scan and the mean of the aligned previous 10 A-scans. The alignment seems to be a compromise between aligning the RPE and the hyperreflective top of the RNFL.

5.2 Diffusion Methods

The physical phenomenon diffusion is the spontaneous spreading of for instance heat or particles. It is a consequence of the continuity equation, which ensures the conservation of a property. If used in images the average intensity is conserved, but the distribution of pixelvalues will change, depending on the type of diffusion.

As mentioned in section 2.1 about speckle, several postprocessing methods to reduce the effect of speckle in OCT images have been used. In the following three types of diffusion filtering will be tested on the OCT retinal image with a macular hole shown in figure 5.2.

There are two different goals with the filtering. Beside the obvious one of producing a more accurate and easily interpretable result it is also to produce an image that is easily registered when looking at several images of same area. This is important if we want to reduce noise by averaging. This will be investigated in chapter 6. The accuracy of the different methods will be tested quantitatively by comparing the result against an estimate of ground truth, achieved by averaging several noisy images.

5.2.1 Linear Isotropic Diffusion

The simplest type of diffusion is the linear isotropic. Isotropic meaning that the diffusion is independent of direction. It is described on an image $I_0(x, y)$ by a set of derived images $I(x, y, t)$ by the following equation

$$\frac{\partial I}{\partial t} = c\Delta I, \qquad I|_{t=0} = I_0, \qquad 0 < c \in \mathbb{R} \tag{5.4}$$

where Δ is the laplacian. This may, as pointed out in [29], equivalently be viewed as convolving the original image with a Gaussian kernel $G(\sigma)$ of standard deviation of $\sqrt{2t}$, when $c = 1$, so

$$I(x, y, t) = I_0(x, y) * G(\sqrt{2t}). \tag{5.5}$$

One problem when using linear diffusion is that edges are smoothed along the flow over time. As a consequence the zero crossings of the second derivative, which indicate the locations of edges, vary over time. To overcome this problem Perona and Malik proposed a nonlinear diffusion in [30], called anisotropic diffusion, described by

$$\frac{\partial I}{\partial t} = \nabla \cdot (c(|\nabla I|)\nabla I), \qquad c(\cdot) > 0 \tag{5.6}$$

where ∇ is the gradient operator and the diffusion coefficient c is a decreasing function of the gradient. This leads to a relation between the magnitude of the gradient in a given direction and the amount of smoothing done in this direction, ie. it is anisotropic. The greater the magnitude, the less smoothing is done. This diffusion method has not been tested, since the next presented is an extension and seems to have proven its superiority.

5.2.2 Complex Diffusion

The process of diffusion can be generalized further. Gilboa suggested in [29] a way to do complex diffusion. It is similar to anisotropic diffusion, except that I and c are complex, and the diffusion filter is a function of the imaginary part of I.

$$\frac{\partial I}{\partial t} = \nabla \cdot (c(Im(I))\nabla I). \tag{5.7}$$

For an appropriate choice of the function $c(\cdot)$ this will lead to the imaginary part of I converging towards the laplacian of the image. The reason for choosing the laplacian instead of the gradient as an indication of the location of the edges, is that the laplacian is high at the endpoints of a ramp, and not within, which leads to better smoothing within the ramp. One problem with the anisotropic diffusion has been a staircasing effect, ie. the edges are enhanced. This is overcome with the choice of the diffusion coefficient being a function of the laplacian in complex diffusion. For this reason it is also called Ramp Preserving, when

$$c(Im(I)) = \frac{e^{i\theta}}{1 + (\frac{Im(I)}{k\theta})^2}, \tag{5.8}$$

where k is a threshold parameter, and θ is the phase angle that should be small, $\theta << 1$. It can be shown that for a small θ an approximation to the diffusion will be

$$Re(\frac{\partial I}{\partial t}) \approx Re(\Delta I), \tag{5.9}$$

$$Im(\frac{\partial I}{\partial t}) \approx Im(\Delta I) + \theta Re(\Delta I), \tag{5.10}$$

which means that the real part of I is dominated by a linear diffusion process and the imaginary part is affected by both the imaginary and real part. Since I is real at $t = 0$, it will approximate the laplacian of the image, and $Im(\Delta I)$ can be regarded as a smoothing term. The derivation and implementation can be seen in [29]. The article also has visual comparisons between the effect of the anisotropic diffusion and complex diffusion on different images.

These basic diffusion equations can of course be extended with other well known schemes. Anisotropic diffusion has been combined with a multiscale laplacian pyramid approach, which in [31] has been demonstrated to reduce speckle better than a wavelet domain speckle reduction technique reported in [32].

5.2.3 Coherence Enhancing Diffusion

Another way to extend the diffusion equations is to incorporate coherence information, ie. the local structure and direction of the local structure. This can be done by replacing the diffusion coefficient by an appropriate 2×2 diffusion tensor D as is done by Weickert in [33]

$$\frac{\partial I}{\partial t} = \nabla \cdot (D(I)\nabla I). \tag{5.11}$$

Before describing how the diffusion tensor is defined, the structure tensor needs to be introduced. Let $I_\sigma = I_0 * G(\sigma)$ be a slightly smoothed version of the original image, the symmetric and positive semi-definite matrix

$$J_\rho(\nabla I_\sigma) = (\nabla I_\sigma \nabla I_\sigma^T) * G(\rho) \tag{5.12}$$

is called the structure tensor. The first element is the second order derivative in the first direction, the second and third element is $\frac{\partial^2 I}{\partial x \partial y}$ and the fourth element is the second order derivative in the second direction. The reason for the first convolution is to reduce noise before determining the derivatives. The second convolution, which should have a larger standard deviation than the first, is for finding the local orientation at an appropriate scale.

The two eigenvalues $\lambda_1 > \lambda_2$ of the structure tensor are used as indicators of local structure in the image with respect to the eigendirections e_1 and e_2. The direction e_1 is where there is the greatest local variation. That means the direction e_2 can be used as an estimate of the direction with the greatest coherence, and the value $\kappa = (\lambda_1 - \lambda_2)^2$ is used as a measure of how well defined the direction of the coherence is.

The diffusion tensor D which indicates the strength and orientation of the diffusion can now be defined

$$D = \begin{pmatrix} \alpha & 0 \\ 0 & \alpha \end{pmatrix} \quad \text{if } \lambda_1 = \lambda_2 \text{ else} \tag{5.13}$$

$$D = (e_1, e_2) \begin{pmatrix} \alpha & 0 \\ 0 & \alpha + (1-\alpha)e^{-\kappa_0/\kappa} \end{pmatrix} (e_1, e_2)^T, \tag{5.14}$$

where α is a small positive constant and κ_0 is a cutoff parameter for well defined orientation. It is noted that the diffusion tensor has the same eigenvectors as the structure tensor. Almost no diffusion is done in the direction with a high structural variation. α was introduced, besides theoretical reasons, to make sure that the diffusion process never stops even if the structure becomes isotropic.

5.2.4 Test Parameters

In the following complex diffusion and coherence enhancing diffusion will be evaluated against linear isotropic diffusion, on an appropriate test image and all methods will be tested on a noisy OCT image of the retina. To evaluate the performance quantitatively the signal to noise ratio (SNR) is calculated. It will be defined as

$$SNR = \sum_{i=1}^{N}\sum_{j=1}^{M} I(i,j)^2 \Bigg/ \sum_{i=1}^{N}\sum_{j=1}^{M} (\widehat{I}(i,j) - I(i,j))^2, \tag{5.15}$$

where I and $\widehat{I}$ are the original image and the denoised image respectively. We are interested in suppressing speckle noise but still preserve edges of the original image, since boundaries in OCT images often are the regions of interest. In order to evaluate the methods abilities to preserve edges, a parameter β originally defined in [34] is also determined. The closer β is to unity, the better the edges are preserved

$$\beta = \frac{\Gamma(\Delta I - \overline{\Delta I}, \Delta \widehat{I} - \overline{\Delta \widehat{I}})}{\sqrt{\Gamma(\Delta I - \overline{\Delta I}, \Delta I - \overline{\Delta I}) \cdot \Gamma(\Delta \widehat{I} - \overline{\Delta \widehat{I}}, \Delta \widehat{I} - \overline{\Delta \widehat{I}})}}, \tag{5.16}$$

where ΔI and $\Delta \widehat{I}$ are the highpass filtered versions of I and $\widehat{I}$ respectively, obtained with a 3×3 standard approximation of the laplacian operator, and

$$\Gamma(I_a, I_b) = \sum_{i=1}^{N} \sum_{j=1}^{M} I_a(i,j) \cdot I_b(i,j). \tag{5.17}$$

5.2.5 Test of Complex Diffusion

To test the implementation of the complex diffusion method, a test image of 100×100 has been constructed. It has linearly varying pixelvalues in one direction and quadratic in the other direction, $I_{test}(r,c) = r + \frac{c^2}{90}$, except for the 10 border pixels where $I_{test} = 10$ and a diagonal line where $I_{test} = 180$. The test image is shown in figure 5.3(a). From this test image a noisy one has been constructed as

$$I_{noisy} = (1 + n_m) I_{test} + n_a, \tag{5.18}$$

where n_m are realizations of a gaussian with zero mean and a standard deviation of 0.1 and n_a are realizations of a gaussian with zero mean and a standard deviation of 10. The noisy test image is shown in figure 5.3(b). Linear diffusion and complex diffusion have been tested with different parameters, but only one of the best results achieved with each method is shown in figures 5.3(c) and 5.3(d). With linear diffusion the edges degrade rapidly, so a compromise between noise reduction and edge preservation is achieved with a standard deviation of $\sigma = 1$. The method yields a SNR of 53.9 and $\beta = 0.41$. A satisfactory result is achieved with complex diffusion when $k = 0.05$ and after 20 iterations. For this instance and every instance run on OCT images, θ in equation 5.8 is kept at $\frac{\pi}{30}$ and the incremental timesteps 0.2. This gives a SNR of 69.2 and $\beta = 0.74$.

The result achieved with complex diffusion is significantly better than what linear diffusion produces. The edges are sharper, and the slowly varying regions are better smoothed. This is reflected in the β and SNR, where both are significantly better for complex diffusion.

The only thing not completely satisfactory is the smoothing and lack of lightness preservation of the white diagonal line. The reason the method has problems

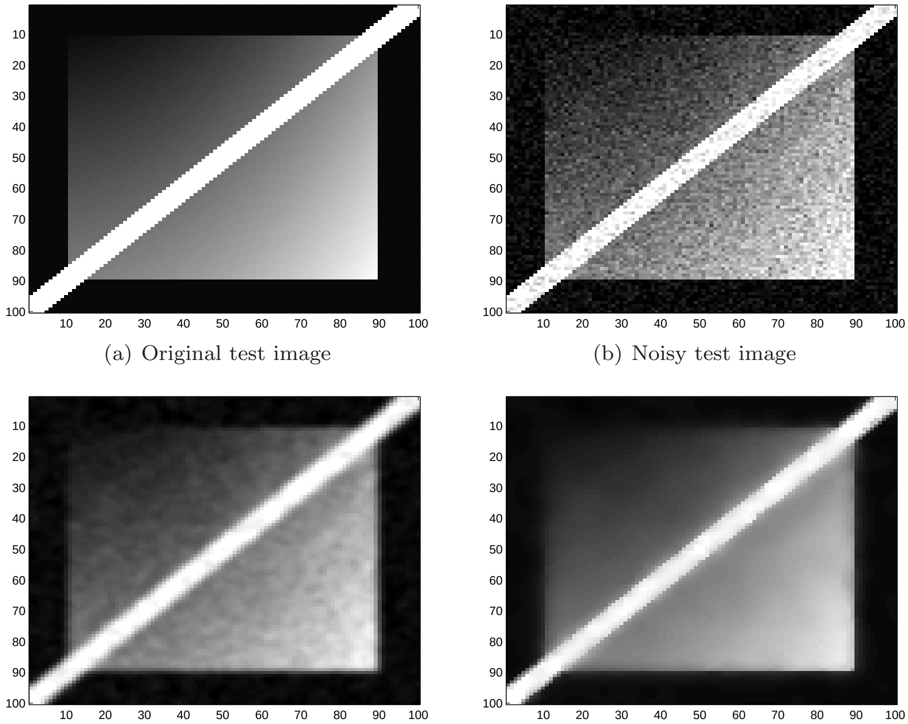

(a) Original test image

(b) Noisy test image

(c) Result of linear diffusion

(d) Result of complex diffusion

Figure 5.3: The noisy test image is denoised by linear diffusion with $\sigma = 1$ and complex diffusion with $k = 0.05$ and 20 iterations. Complex diffusion clearly does a better job at preserving edges and smoothing the continuous areas. The only thing not completely satisfactory is the smoothing and lack of lightness preservation of the white diagonal line.

right here is probably due to the multiplicative noise and high pixelvalues that gives a significant noise level when looking at the absolute values of the noise.

5.2.6 Test of Coherence Enhancing Diffusion

If coherence enhancing diffusion is tested on the same image, the strength of the method would not be apparent, so an image of a fingerprint is used as a test image instead. This is because enhancement of fingerprints is one area the method has proven its worth. To calculate the two parameters, SNR and β, it is necessary to have a golden standard to measure up against. Since such an image is not available in the case of the fingerprint image, these values are not calculated. The original fingerprint image is shown in figure 5.4(a), the result achieved by linear isotropic diffusion and coherence enhancing diffusion can be seen in figure 5.4(c) and 5.4(d), respectively.

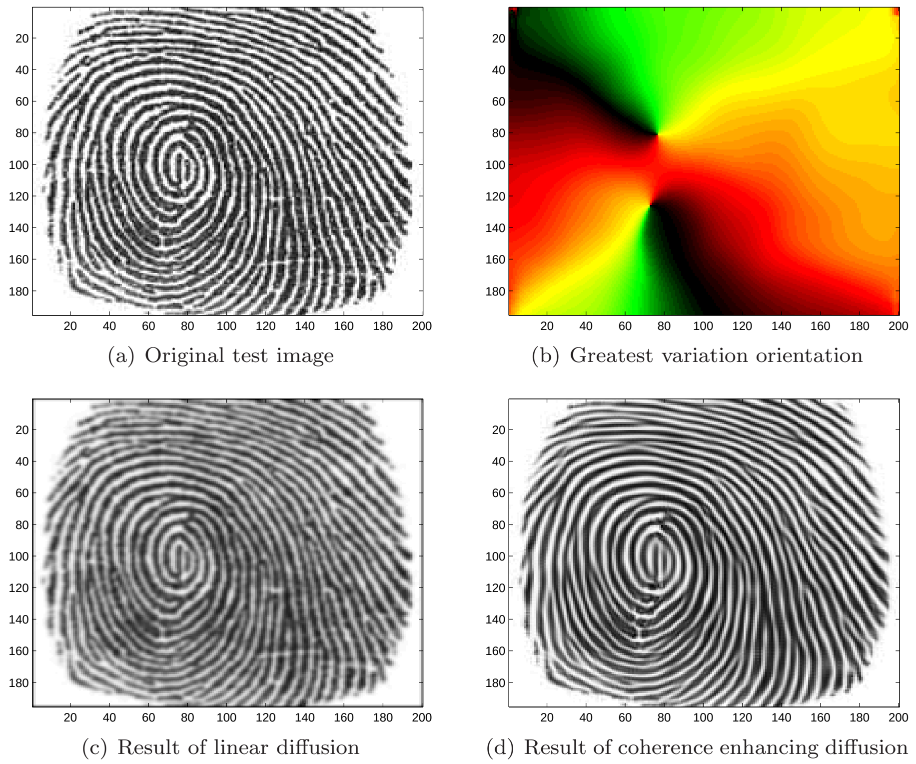

(a) Original test image

(b) Greatest variation orientation

(c) Result of linear diffusion

(d) Result of coherence enhancing diffusion

Figure 5.4: The fingerprint image in (a) is denoised by linear diffusion with $\sigma = 1$ and coherence enhancing diffusion with no σ, $\rho = 10$, 10 iterations and κ_0 set to the 25% quantile of κ. The result really shows how coherence enhancing diffusion smoothes discontinuities across the coherence direction, but maintains the sharpness between light and dark lines. The orientation of the eigenvector corresponding to the largest eigenvalue of the structure tensor is shown in (b). Red is horizontal and green is vertical and the diagonals are black and yellow.

By visual inspecting figure 5.4(d), it can really be seen how well coherence enhancing diffusion works in the case of images with a high directional coherence. Discontinuities in the coherence direction are smoothed, but the sharpness between light and dark lines perpendicular to the coherence direction is maintained. The orientation of the eigenvector corresponding to the largest eigenvalue of the structure tensor is shown in 5.4(b). Red is horizontal and green is vertical and the diagonals are black and yellow. The parameters used in this case is no σ, $\rho = 10$, 10 iterations of timestep 1 and κ_0 set to the 25% quantile of κ. For this test and all applications on OCT images, $\alpha = 0.001$ and timesteps are set to 1.

5.2.7 Test on OCT image

In the case of OCT images, a speckle free image is needed as a golden standard. Since such an image is not available, a result of averaging 11 registered images of the same region of a retina with a macular hole is used. The procedure will be explained in section 6 and the result is shown in figure 6.9 on page 55. It is obvious that noise due to speckle is significantly reduced compared with figure 5.2.

The two parameters SNR and β have been calculated for four values of the standard deviation for the linear isotropic diffusion. For higher standard deviations both of the parameters decrease in value, and the denoised image deteriorates. The parameters have been calculated for four different values of complex diffusion and two values for coherence enhancing diffusion.

Method		SNR	β
No filtering		8.056	0.496
LID	$\sigma = 1$	19.972	0.293
	$\sigma = 2$	27.607	0.091
	$\sigma = 3$	31.495	0.045
	$\sigma = 4$	32.475	0.030
CD	$k = 0.05,\ \text{iterations} = 4$	20.794	0.367
	$k = 0.5,\ \text{iterations} = 9$	27.756	0.095
	$k = 1,\ \text{iterations} = 20$	31.707	0.052
	$k = 1,\ \text{iterations} = 30$	32.638	0.042
CED	$\sigma = 1,\ \rho = 10,\ \kappa_0 = 10\%$ quant, ite $= 30$	22.044	0.253
	$\sigma = 3,\ \rho = 10,\ \kappa_0 = 10\%$ quant, ite $= 30$	24.532	0.205

Table 5.1: Values of SNR and β achieved with different parameters for linear isotropic diffusion, complex diffusion and coherence enhancing diffusion.

The denoised images when using complex diffusion can be seen on figure 5.5. The denoised images achieved with linear isotropic diffusion are not shown, since they are visually similar to what is achieved with complex diffusion, in the way

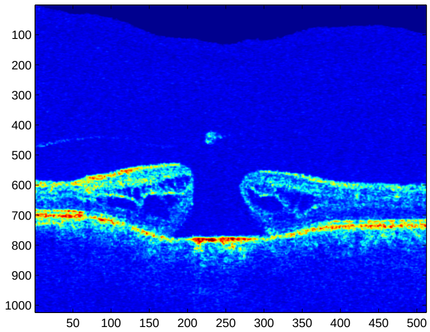

(a) $k = 0.05$, iterations $= 4$

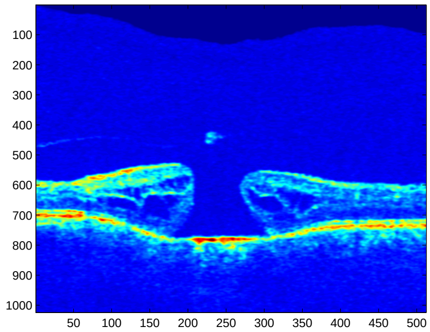

(b) $k = 0.5$, iterations $= 10$

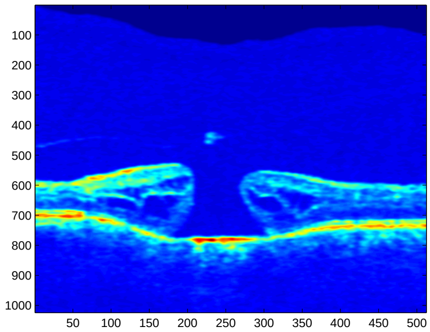

(c) $k = 1$, iterations $= 25$

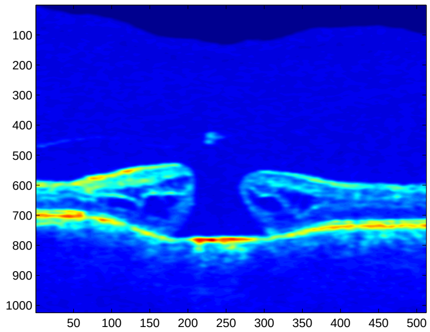

(d) $k = 1$, iterations $= 40$

Figure 5.5: Denoised images achieved with complex diffusion with four different sets of parameters. The results are very similar to what is achievable with linear isotropic diffusion.

that the denoised image achieved with $\sigma = 1$ is similar to 5.5(a), $\sigma = 2$ is similar to 5.5(b), $\sigma = 3$ is similar to 5.5(c) and $\sigma = 4$ is similar to 5.5(d). As seen in table 5.1 it is possible to produce denoised images with complex diffusion that have higher SNR and β for every integer value of the standard deviation with linear isotropic diffusion. But it is also noted that the improvement is minimal, especially when visually inspecting the results. The resulting images achieved with complex diffusion are very similar to what can be achieved with linear isotropic diffusion. This may be because gaps in hyperreflective areas that diffusion should flow across have the same properties as actual hyporeflective areas that should be preserved. If the method can not distinguish between the two, none or both of the two cases will be diffused, and none of the possibilities are satisfactory.

The SNR and β referred to here should be taken with some caution, since as previously mentioned no absolute noise free image was available. But the average image used as ground truth has significantly less noise than one of the originals.

Since two variables have been adjusted instead of one, the implementation of complex diffusion is not as intuitive and easy as linear isotropic diffusion. But in general the higher k and iterations, the more smoothing is done.

When adjusting σ and ρ for the coherence enhancing diffusion method, it is seen that it is very hard for it to pick up the appropriate coherence directions in the noisy OCT image. If σ, controlling the noise suppression smoothing, is increased above three, the borders lying between the edemas, ie. the large non reflecting areas in the retina, tend to be diffused in a wrong direction and hence disappear. This is of course not acceptable. If σ is too small, the estimated coherence direction tend to be vertical for almost the entire image, as is seen in figure 5.6(c). This is because each A-scan is highly correlated locally even as an original image. The SNR values for coherence enhancing diffusion are not impressive. This can partially be because of the sometimes wrongly estimated coherence orientation, but also since very little general smoothing is done in places that are close to isotropic. This leads to a larger SNR than an image with the average value in all the pixels in the isotropic area would.

In figure 5.6(b) a more general horizontal structure is found within the retina, but there are still a few places of an estimated vertical coherence direction, eg. in the top right part of the retina where a structure is created that is questionable. This ability, even though the images created are pleasing to look at is simply not acceptable if an ophthalmologist should use the image for diagnosis. Coherence enhancing diffusion is not tested as a preprocessing step for registration of several images, primarily because of its ability to generate coherent structures that most likely is due to noise.

5.2.8 Conclusion

It is possible to increase SNR by more than a factor of four with the use of diffusion. You might even say that the images looks visually more pleasing for the human eye, but there are limitations when only one image is available. With complex diffusion the results are only marginally better than what is achieved by simple gaussian filtering. With coherence enhancing diffusion, the problem arise that coherence directions can be forced where none exists. When evaluating the three diffusion methods, the conclusion can be drawn that the noise due to speckle is not significantly reduced, without compromising the image in other

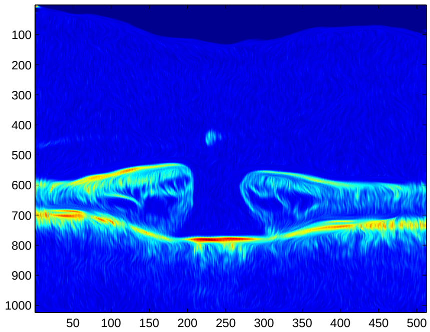

(a) $\sigma = 1$, $\rho = 10$, $\kappa_0 = 50\%$ qua, ite $= 20$

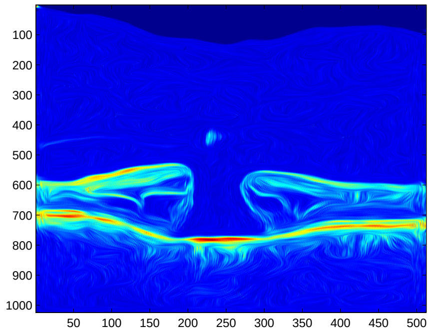

(b) $\sigma = 3$, $\rho = 10$, $\kappa_0 = 50\%$ qua, ite $= 20$

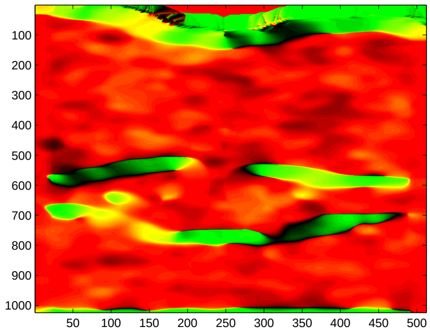

(c) Greatest variation orientation for (a)

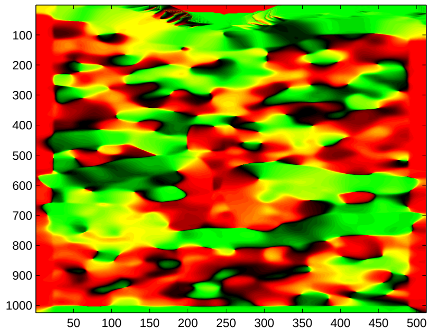

(d) Greates variation orientation for (b)

Figure 5.6: Denoised images achieved with coherence enhancing diffusion with two different sets of parameters are shown in the top two images. The orientation of the eigenvector corresponding to the largest eigenvalue of the structure tensor is shown below each image. With the parameters used in the left case, almost the entire image is estimated to have a vertical coherence. The image is therefore smoothed in this direction. With a larger σ, the structure tensor is estimated from an image where more details have been smoothed, which leads to a higher degree of what would be classified as correct horizontal structure.

ways. One problem is that the speckle plays a dual part in OCT, it is a source of noise, but it is also a carrier of the signal [11], and it is not an easy task to distinguish between the two effects. But if more than one image of a given area is available, they can be combined to reduce the noise without significantly smoothing the signal as shall be seen in the next section.

Chapter 6

Multiple Images

The idea behind averaging several images comes from the fact that uncorrelated noise will be reduced, when several samples are averaged. If a set of signals with uncorrelated gaussian noise is averaged, the standard deviation of the average of the signals is reduced a factor of the square root of the number of signals, compared to one of the original signals. Any correlation between the noise in the signals will reduce this effect. The speckle pattern will not change if no changes in the measuring geometry has happened, but since small movements occur in the eye relative to the apparatus, as was discussed in section 2.1, a large part of the correlation in the speckle is lost. In addition it must not be forgotten that any uncorrelated noise not due to speckle that exists in the image will be reduced significantly by averaging.

6.1 Registration

Averaging of the images will not make any sense before they are registered. The term vertical registration will refer to the process of registering every A-scan in an image to the corresponding A-scan in another image. The shifts that need to be applied to every A-scan need not be the same. By horizontal registration will be meant the process of shifting the entire image horizontally to match with another image. Every row must be shifted the same amount.

When a retinal OCT image is being taken of a retina, the individual focus on a red dot. This makes it easier for the ophthalmologist to set the image plane right through the center of the fovea. Therefore only limited horizontal movement is seen between images taken of an individual without any pathologies. This is not the case when dealing with an individual with a macular hole, since it can be virtually impossible for them to focus on the red dot. This causes an increase in the variation of the location of the image plane. The variation has been illustrated in figure 6.1 that shows the location of three OCT B-scans on the back of the eye. If the desired plane is the one indicated by the red line, then a scan taken at the green location is of limited use, and must be removed from the set, if a significant reduction in noise is wanted. But if the points of interest is located at the center of the red line, the scan located at the blue line can still be of use if it is registered horizontally with respect to the image plane.

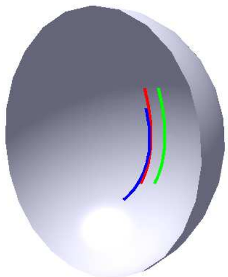

Figure 6.1: Three different located OCT B-scans on the back of the eye. If the desired scan is the one indicated by the red line, then a scan taken at the green location is of limited use. But the scan located at the blue line can still be of use if it is registered horizontally with respect to the image plane.

6.1.1 Vertical Registration

It will now be assumed that all the images in the set are lying on approximately the same plane, ie. we are allowing translations such as the one occurring between the red and blue line, but not red and green line. The process of registering two OCT images will be illustrated with the two images shown in figure 6.2(a) and 6.2(b).

One way to register these two images vertically would be to maximize the cross-correlation for every two corresponding A-scans. This can be illustrated in what will be called energyspace, which for every column is defined as the negative cross-correlation between the corresponding two A-scans in the images a and b

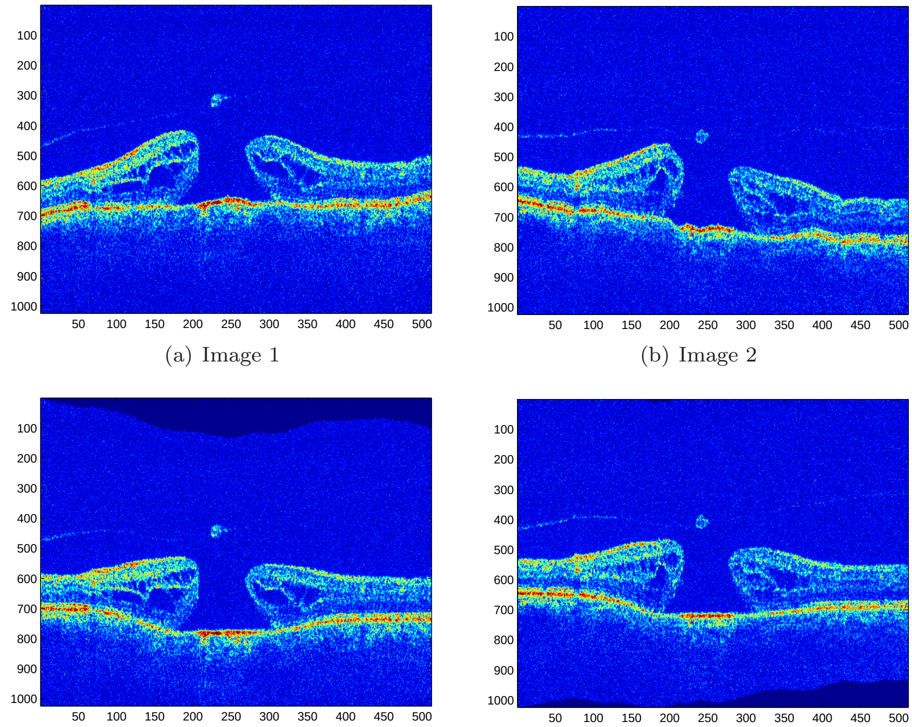

(a) Image 1 (b) Image 2

(c) Aligned image 1 (d) Aligned image 2

Figure 6.2: Image (a) and (b) shows two OCT images taken of the same retinal location. Image (c) and (d) are aligned versions of (a) and (b) respectively.

that needs to be vertically registered. The energyspace is based on the cross-correlation as follows

$$\varepsilon(m, j) = -R_{a_j b_j}(m), \tag{6.1}$$

where ε is calculated for every possible shift $-N < m < N$ and every A-scan $1 \leq j \leq M$.

For the two aligned versions of figure 6.2(a) and 6.2(b) shown in figure 6.2(c) and 6.2(d) the energyspace is shown in figure 6.3(a). A point in energyspace $\varepsilon(r, c)$ correlates negatively with how likely a vertical shift of r pixels between columns c is. At the dark positions there is a large correlation between the two A-scans. The structure of the two images can be recognized in energyspace. The horizontal line at around 1100, which corresponds to a shift of little less than 100 applied to figure 6.2(d), is where the two images line up. The lines above and below is where the RPE in one image line up with the RNFL in the other.

One thing worth noting is that the top and bottom of energyspace is very light, since there is almost no correlation. This means that there is no chance that a solution close to the top or bottom will be found. Using the cross-correlation without normalizing with the size of the overlapping interval has shown to be very robust.

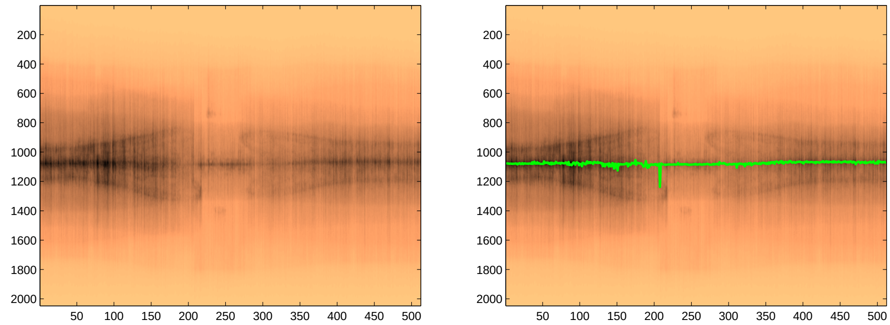

(a) Energyspace (b) Naive solution

Figure 6.3: Image (a) shows energyspace between image 6.2(c) and 6.2(d). It is defined as the negative cross correlation between the corresponding A-scans in the images. In image (b) a naive solution to the vertical registration problem is shown.

The naive implementation of taking the maximum correlation for every A-scan corresponds to the solution shown with the green line in figure 6.3(b). The method is sensitive to noise in two ways. It allows very sudden changes in the shifts from A-scan to A-scan, as can be seen one place in the figure, and it fits to noise on a small scale, which can be seen on the roughness of the green curve. The first problem can be fixed the same way as was done when aligning an image, by only allowing a certain maximum change from A-scan to A-scan, but the other problem can not as easily be fixed.

The solution in energyspace we are interested in, is a solution that is smooth. We want to utilize the horizontal structure that exists in the images and therefore also in energyspace. This can be achieved by finding the shortest path from one side of energyspace to the other with constraints on the curvature.

Since we assume inherent horizontal structure in the images, two images that have not been horizontally registered can still be vertically registered. This will give a first estimate of the vertical shifts that needs to be applied to each column to register the two images.

Finding the shortest path has been implemented with regularized dynamic programming, which was described in chapter 4. It finds the optimal path through energyspace crossing the image from one side to the other, with constraints on the shape of the path. These constraints makes the difference, when comparing the method with standard dynamic programming. As was discussed in chapter 4, the regularized version will have smoother solutions depending on the regularization constant λ and how energyspace is sampled. A result with $\lambda = .1$ and $p = 1$ is shown in figure 6.4(a). The shortest path is smooth and close to horizontal, indicating that the two images have been aligned well.

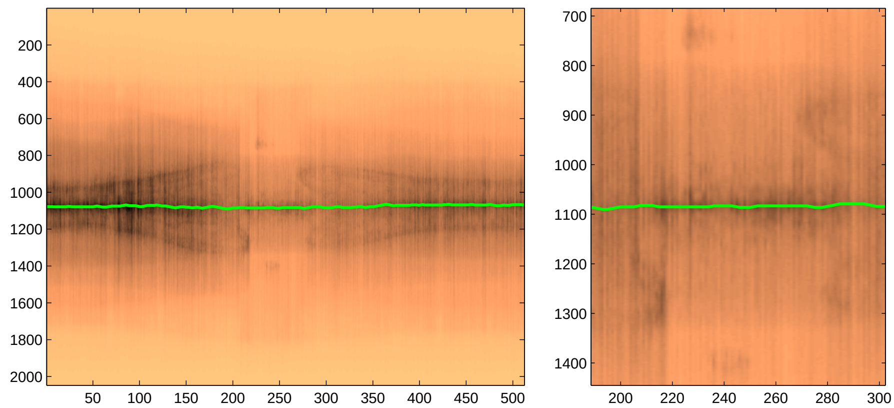

(a) Regularized dynamic programming (b) Zoom of (a)

Figure 6.4: The shortest path achieved with regularized dynamic programming with $\lambda = .1$ and $p = 1$ is shown in (a). The path is smooth and close to horizontal, indicating that the two images have been aligned well. In (b) is shown a closer look at the central part of the path. A significant change in energyspace values column to column can be seen, but the minimum values are close to lined up.

The method has been tested with different values for λ, p and different rescalings, but visual comparison leads to the two initial conclusions that rescaling does not improve the result and there is no reason for increasing p beyond 1. The reason that pixel subdivision does not improve the result may be because the path changes gradually with a low degree of horizontal correlation, which means that no large part of the path with low slope is needed. Horizontal subsampling may be a good idea, if there is a high degree of noise variation column to column. As can be seen in figure 6.4(b), where the central part of the shortest path is shown, there is a high degree of variation in pixelvalues column to column, but the minimum values tend to lie in the same vertical locations. The constraints on the path is enough to eliminate the noise effects. Finally the reason that higher values of p does not improve the result is since the change in direction in the optimal path happens gradually as well, and there is no need for allowing larger slopes.

Registering the image shown in 6.2(d) to the one in 6.2(c), and averaging the two gives the noise reduced image shown in figure 6.5. Compared with the original images, there is a reduction in noise, eg. in the RPE. But there is also a problem. Looking at the sides of the hole, it can be seen that the two images are not horizontally registered. The images must therefore also be horizontally shifted accordingly. The predecessor method to the one described in this chapter, developed by Jørgensen [5] did not incorporate this. This was at the time not necessary, since the pathologies being investigated, had less horizontal variation than patients affected by a macular hole. This larger variation is, as previously mentioned, because it is difficult for the patients to focus on a dot, while the image is taken, since their central vision is impaired.

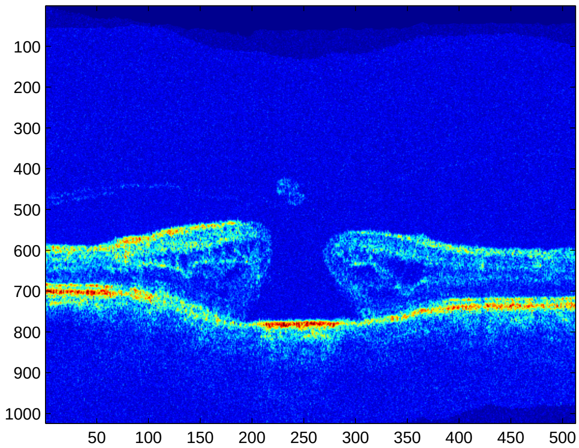

Figure 6.5: Average of the two images 6.2(c) and 6.2(d) after vertical registration. Compared with the original images, there is a significant reduction in noise. But looking at the sides of the hole, it can be seen that the two images are not horizontally registered.

6.1.2 Horizontal Registration

As previously described in the first part of section 6.1 the images are normally close to being horizontally registered, unless a pathology affects the visual acuity in the macular area, which is the case with macular holes. Since we assume no movements have happened except in the longitudinal direction during the recording of the image, only one horizontal shift needs to be estimated, this is again done by maximizing cross-correlation, this time horizontal. The images are transposed, this way the indices of a^T and b^T corresponds to rows and not columns in the original images.

$$\max_m \sum_{r=1}^{N} R_{a_r^T b_r^T}(m) \tag{6.2}$$

Again the cross-correlation is robust, but a larger shift means that parts of the images will not be included. This was not a problem when registering vertically, since almost no signal existed at the top and bottom. If this is not corrected, the optimal shift will be biased toward no shifts, so to overcome this, the unbiased estimator is used. This means a factor $\frac{1}{M-|m|}$ is multiplied to equation 6.2. The unbiased estimator can be unstable for extreme shifts, so m is only searched in the interval $-100 < m < 100$, which is more than sufficient for any realistic case. The average of the two images 6.2(c) and 6.2(d) after vertical and horizontal registration is shown in 6.6. Image 6.2(d) is shifted 9 pixels to the left. The relative limited shift makes a significant change to the resultant image. The edges of the hole are much sharper and the hyperreflective point to the far left can be localized. The roof of the foveal cyst suspended above the hole, is not vertically aligned though. This can not be achieved with the current method, since the RPE in the same A-scans are aligned. To improve this a stretching of the A-scans is needed. This will not be investigated, since there is no physical evidence of this happening. The reason for the not aligned roof is probably due to actual movement of the roof or the scans have not been taken in the exact same plane as shown in figure 6.1.

6.2 Procedure

It has now been shown how two images can be registered. In this section the procedure is extended to any number of images by iteration. The entire procedure is for clarity shown in the flowchart in figure 6.7.

An estimate of the registration of image one and two can be determined as described in the previous section. The mean of the two registered images is used as a template for image three. The procedure is repeated, until all images have been included. This is what is marked with a box as the first iteration. The second iteration consists of re-registering every image to the template, without updating the template. The reason for this is that each image is registered according to the same template, which means the correlations found in energyspace can be compared between the images. This is used when the final noise

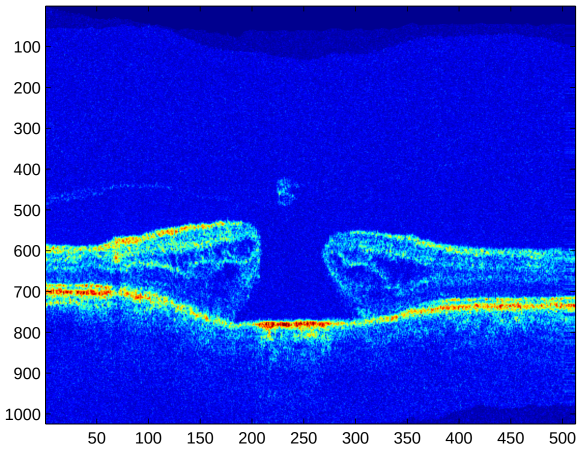

Figure 6.6: Average of the two images 6.2(c) and 6.2(d) after vertical and horizontal registration. The relative limited shift makes a significant change to the resultant image.

reduced image is being produced, which is the weighted mean of all the original images registered according to the second iteration. What weights to use will be looked at later on.

6.2.1 Test of Procedure

When the text in a box in flowchart 6.7 is followed by an asterisk, it means that a test at this point has been performed. The procedure has been tested in three ways.

- Different types of prefilters

  Since one of the conclusions in section 5.2 was that linear and complex diffusion produced very similar results, it is of minor importance which one is being tested. Because of the ease of interpretation and the speed of the method, linear diffusion is tested. Coherence enhancing diffusion is not tested as a prefilter, since it had an ability to create structures that may not be present, which could lead to a less correct registration. The method have been tested without prefiltering, and with a gaussian filter with standard deviation equals to 1, 2 and 3.

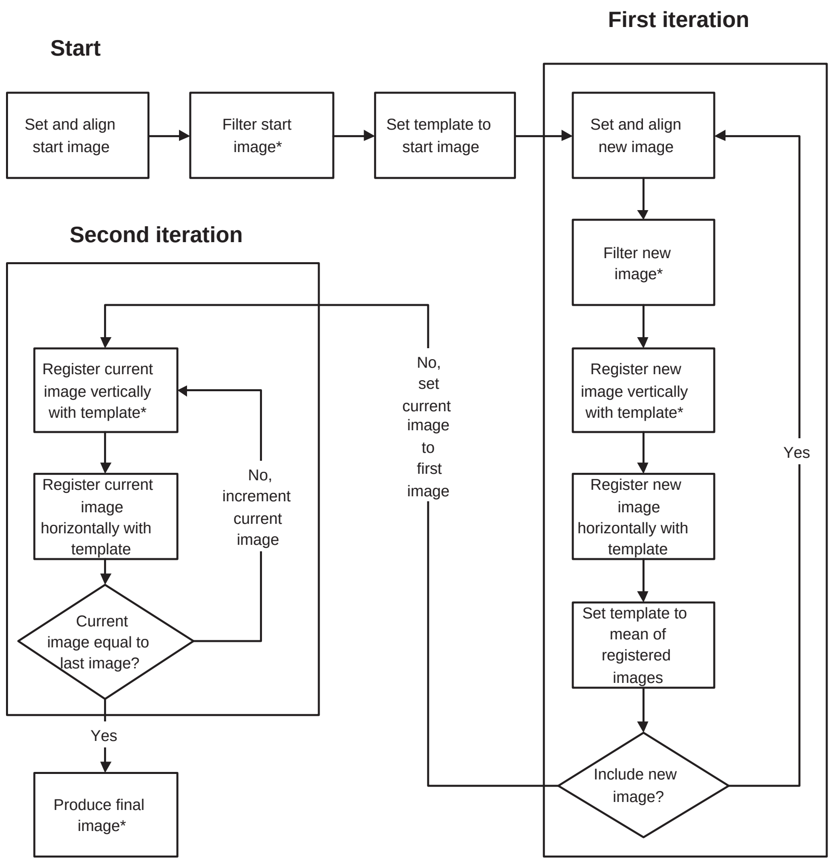

Figure 6.7: Flowchart of the procedure of registering several images and producing a final noise reduced result.

- Different shape constraint in regularized dynamic programming

  The location of the shortest path in energyspace depended on the regularization constant λ. The method has therefore been tested for λ equal to 0, 0.1 and 0.5, where 0 is equal to dynamic programming without regularization. The method has also been tested with a different horizontal scale factor q, where the horizontal resolution has been reduced a factor of 2.

- Different weights when producing resultant image

  When the resulting image is produced, each column in the images can be weighted according to how much it correlates with the template. These values are known from the values in energyspace the shortest path crosses. The method has been tested with no weights, ie. taking the mean of the registered images, weights equal to the correlation with the template and the squared correlation.

To reduce the number of runs performed during the test, a standard setting has been used, and only one parameter has been changed at a time. The standard setting is using no prefilter, $\lambda = 0.1$ with no change in resolution and averaging the registered images, ie. no final weights.

The test has been performed on two data sets. One is a set consisting of 18 image with a macular hole. The image used throughout this chapter comes from this set. Only 11 of the images have been used in the test. All of these images have a visible roof of the foveal cyst suspended in front of the macular hole, and this is taken as an indication of they are lying in the same plane. The other set consists of 14 images, where all of them were used. This image set is taken 5 weeks after operation, and looks to a layman similar to retinal OCT images without any pathologies. The images have a much greater variation in the horizontal shift than would be the case for images without pathologies though.

To evaluate the methods, a contrast to noise ratio (CNR), inspired by [35], of a part of the final images produced are calculated, and the images are visually inspected.

If two adjacent classes a and b are outlined in an image, one can define the contrast C as

$$C = (\bar{I}_a - \bar{I}_b)^2, \tag{6.3}$$

where $\bar{I}_k$ refers to the mean of pixels in class k. A common squared standard deviation for the two classes can be expressed as

$$\begin{aligned} \sigma_0^2 &= \frac{1}{N_0} \sum (I_0 - \bar{I}_0)^2 & (6.4) \\ &= \frac{1}{2N_a} \sum (I_a - \bar{I}_s)^2 + \frac{1}{2N_b} \sum (I_b - \bar{I}_b)^2 & (6.5) \\ &= \frac{1}{2}\sigma_a^2 + \frac{1}{2}\sigma_b^2, & (6.6) \end{aligned}$$

if the two classes are of the same size, which is the case in the following. The determined CNR is therefore

$$CNR = \frac{2(\bar{I}_a - \bar{I}_b)^2}{\sigma_a^2 + \sigma_b^2}. \tag{6.7}$$

If two noisy classes are found this will affect the denominator, and the CNR will therefore decrease. If too much smoothing has been done, the transition between the two classes is less sharp. This will affect both the numerator and the denominator such that the CNR decreases. It is therefore expected that this measure will be able to reward both less noise and a sharp transition.

The two adjacent classes that have been used is the RNFL, and the vitreous body lying above. The transition has been determined with regularized dynamic programming. The two classes are sampled from 80 columns, 10 pixels for each column. The used part of the macular hole image can be seen in figure 6.8(a) and the used part from the postoperative macular hole can be seen in figure 6.8(b). The images shown are for the standard setting. The transition lines between the two classes are shown on both images.

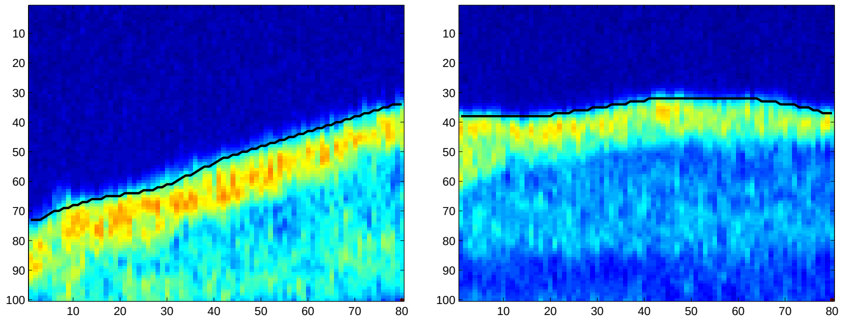

(a) Macular hole (b) Postoperative macular hole

Figure 6.8: The resulting images of the macular hole has had the part shown in (a) cut out. The border has been found and the 10 pixels lying above and below the line have been used as the two classes in the CNR. The same has been done for the postoperative macular hole in (b).

The CNR is listed in table 6.1. Producing these values, it was seen that the values are susceptible to small changes in the location of the transition line. When interpreting the numbers, care must therefore be taken not to over interpret minor variation in CNR.

When visually inspecting the two sets, it is very hard to see a significant change. But one conclusion can be made upon closer inspection. This is that a greater prefiltering leads to a worse horizontal registration. This makes sense, since upon filtering the edges wash away, and minute changes have less effect.

When looking at the CNR values, the first conclusion that can be made is that some prefiltering leads to a higher CNR. It seems as if a standard deviation of

Method	CNR_{set1}	CNR_{set2}
standard settings (ss)	11.3	16.8
ss with $\sigma = 1$	12.1	17.1
ss with $\sigma = 2$	12.6	17.8
ss with $\sigma = 3$	12.1	17.3
ss with $\lambda = 0$	11.4	19.1
ss with $\lambda = 0.5$	10.3	17.3
ss with $\lambda = 0$ and $q = 2$	12.2	16.7
ss with $\lambda = 0.1$ and $q = 2$	11.5	18.6
ss with $\lambda = 0.5$ and $q = 2$	10.6	18.0
ss with proportional weights	11.3	16.6
ss with squared weights	11.3	16.2

Table 6.1: CNR is listed for the 11 test cases, for each of the two test sets.

$\sigma = 3$ reduces CNR, though, so a value less than 3 is optimal, as expected from the visual inspection.

Using a high value for the regularization constant λ seems to constrain the final shifts in such a way that CNR decreases. It may even seem as if no regularization is needed. But a small regularization maintains the horizontal correlation in the produced image, so it is preferred. It is hard to say anything conclusive about reducing the resolution before finding the optimal path. It does not seem to make a significant change. The same seems to be the case for weighting the columns according to their correlation with the template, when the final result is produced. There is a slight tendency that the more the columns are weighted the lower the CNR.

In general the procedure is very robust to noise in the original images. This can be seen since no prefiltering is necessary. It is can also be seen from the fact that the regularization parameter λ, does not have to be included to achieve a satisfactory result.

The final result where a gaussian prefilter with $\sigma = 1$ has been used, $\lambda = 0.1$ and the columns have not been weighted are shown in 6.9 and 6.10(b). The noise reduction is very significant when compared to one of the original images, as shown in 5.2 and 6.10(a). A few details worth mentioning are the layers in general are easily discernable compared to the original images, where some of the layers are not even perceived. There is also a hyperreflective area on the left side of the RNFL in figure 6.10(b) that might not have been noticed in the original image. This area would disappear if no horizontal registration had been done.

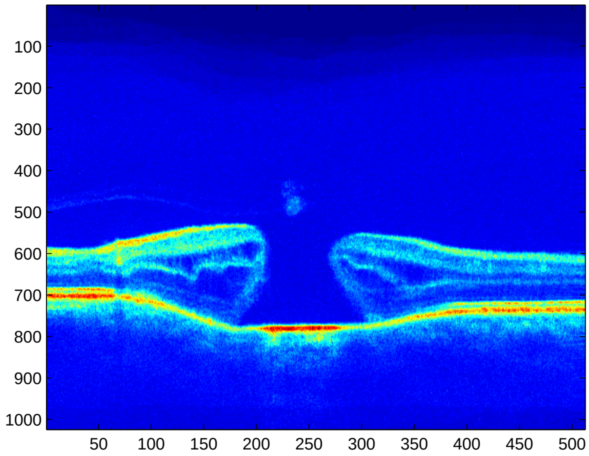

Figure 6.9: Averaging of 11 images a gaussian prefilter with $\sigma = 1$ has been used, $\lambda = 0.1$ and the columns have not been weighted. The noise reduction is very significant when compared to one of the original images, as shown in 5.2.

6.2.2 Noise and Number of Images

Taking an OCT image does not take many seconds, but to minimize the time spend by the ophthalmologist and patient on taking images, it is relevant to know how many images is needed to produce a satisfactory result, or if the improvement stops as a function of images included. For this reason, the resulting image will now be investigated as a function of number of OCT images used.

The CNR is also used to evaluate the final images. The same parts of the images are used as shown in 6.8. For illustration purposes the parts have been shown for the original OCT images in figure 6.11.

The CNR values are listed in table 6.2. As expected there is a significant increase in CNR when including the first couple of images. This effect seem to wear off when 8-10 images have been included. The same is seen visually. Average images for different number of images included can be seen in figure 6.12. It is even hard to see the improvements from 6.12(e) to 6.12(f), where seven and nine images are used respectively. The same variation in CNR, as was previously discussed, due to minor changes in the location of the transition line is also present in these values. Too much should therefore not be put in a variation of one or so in CNR.

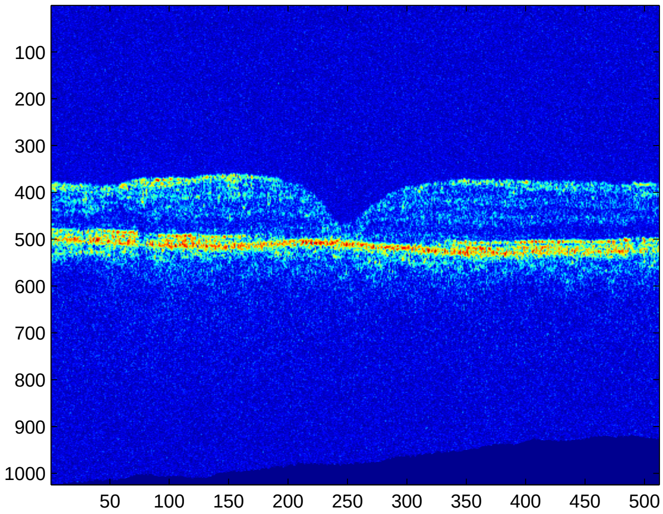

(a) Original image

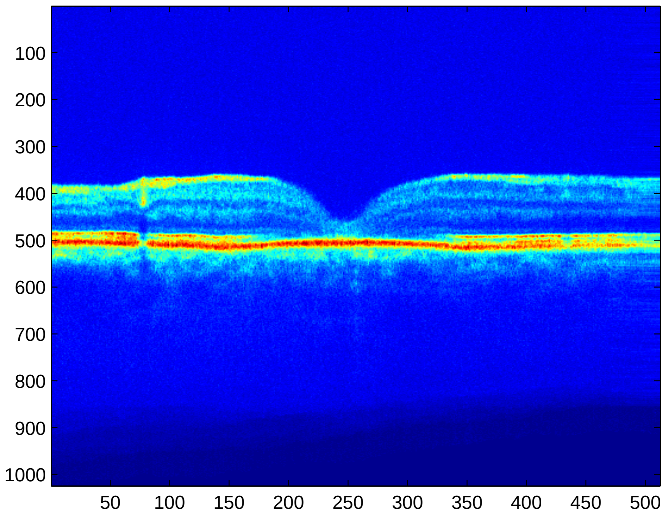

(b) Average image

Figure 6.10: An aligned OCT image of an individual who previously had a macular hole is shown in (a). Image is taken 5 weeks after operation. Average of 14 images with $\sigma = 1$, $\lambda = 0.1$ and no weights is shown in (b). The noise reduction is very significant. The hyperreflective area on the left side of the RNFL is an example of a detail almost not noticeable in the original images.

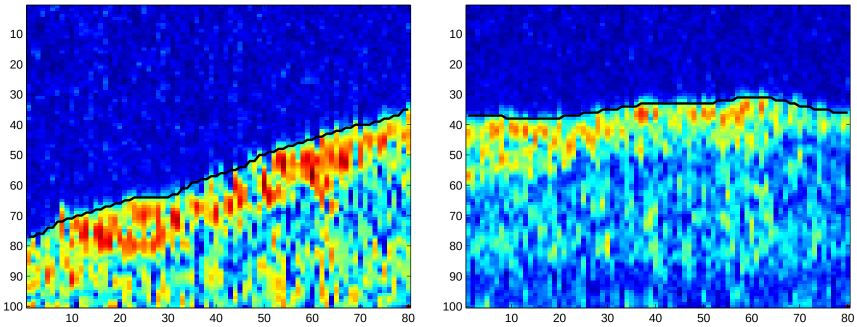

(a) Macular hole (b) Postoperative macular hole

Figure 6.11: The part of all the resulting images of the macular hole has had the part shown in (a) cut out. The border has been found and the 10 pixels lying above and below the line have been used as the two classes in the CNR measure. The same has been done for the postoperative macular hole in (b).

Nr. of images used	CNR_{set1}	CNR_{set2}
1	4.7	4.2
2	5.8	6.9
3	7.1	10.6
4	8.9	11.7
5	8.0	13.4
6	10.0	16.0
7	10.0	16.8
8	10.7	16.8
9	11.4	17.5
10	11.4	16.4
11	12.1	16.5
12		17.3
13		17.9
14		17.8

Table 6.2: CNR is listed as a function of images used to produce the average image, for each of the two test sets.

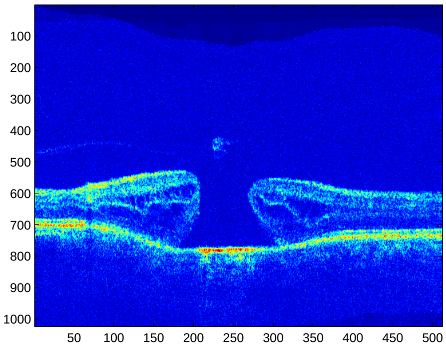

(a) Average of 2 images

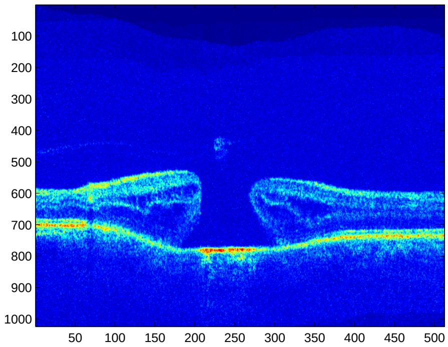

(b) Average of 3 images

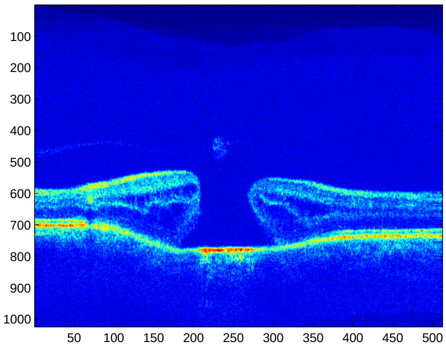

(c) Average of 4 images

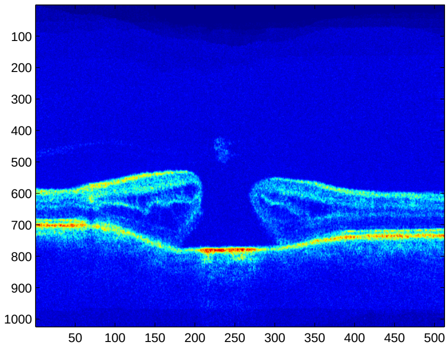

(d) Average of 5 images

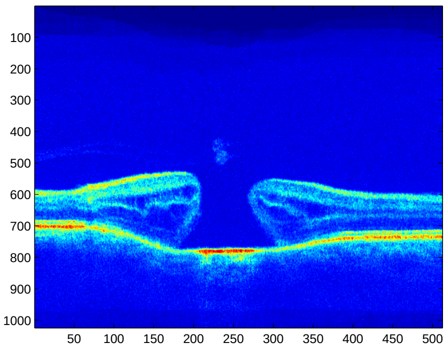

(e) Average of 7 images

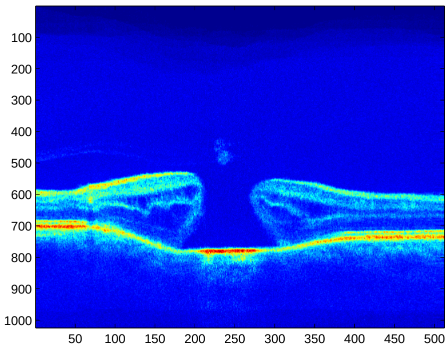

(f) Average of 9 images

Figure 6.12: The resulting image is shown as a function of images included. The original image can be seen in figure 5.2 and the one produced when 11 images are used in figure 6.9.

6.3 Conclusion

A procedure to register a set of OCT images taken of the same retinal location has been presented. The images are not horizontally aligned, which would normally be the case for retinal OCT images. This is because the macular area is affected by the pathology, and the patient can not focus on the red dot in the scanner, which leads to horizontal shifts. This off course also leads to shifts perpendicular to the image plane, but if this shift is too significant, the images must be removed from the dataset before the procedure is started.

The procedure includes one image at a time. First vertical then horizontal registration is done. Due to the horizontal structure it can be assumed that a fairly close to optimal vertical registration is done, although the images are not horizontally registered. Once every image is included, the registration is repeated to the average image, to fine tune the result. The resulting image has significantly less noise than the originals, and details can often be seen that are not apparent in the originals.

Another way to visualize the significant reduction in noise is by looking at a single A-scan as is done in figure 6.13. The red line through the two images indicate the location of the A-scans that are shown to the right. It seems as if the noise reduction has come at the prize of decreasing the height of the peaks. This is a consequence of the averaging of several noisy signals, but it is not known what the true peak heights are, or whether they are closer to the resulting peaks. The original peak values could be due to multiplicative noise.

The procedure has been tested with different parameter settings. The resulting images have been inspected visually, and by measuring the CNR for a part of the RNFL and above lying vitreous body. The resulting images produced with different parameters were very similar, so was the CNR values.

When calculating CNR as a function of number of images averaged, the improvement is significant up to about 8-10. If minute details needs to be investigated, taking more images and discarding any that seems to be shifted perpendicular to the image plane is recommendable.

The method with the chosen parametrical settings have been tested on 7 sets of images. From each set between 10 and 13 images have been used to produce the final image. One set is of a healthy individual, 3 sets have a preoperative macular hole and 3 sets are postoperative images of a macular hole. In Appendix B one initial image, an aligned version of the initial image and the average image for each set are shown. When looking at the resulting images it seems as if all of the used images have been registered correctly both horizontally as vertically.

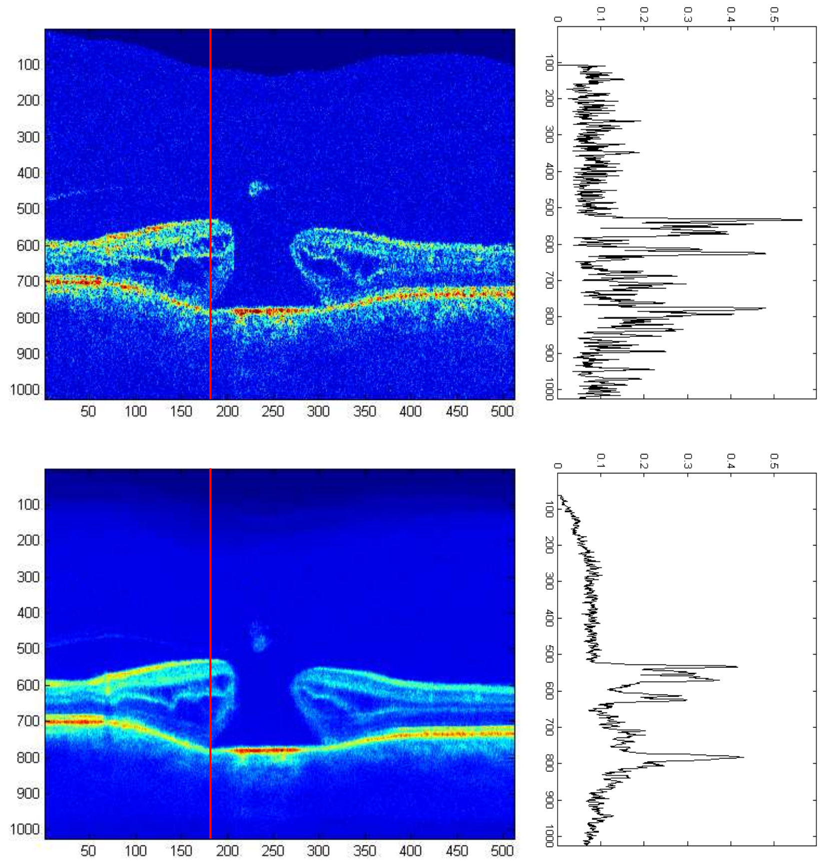

Figure 6.13: Single A-scans are shown to the right of the corresponding image. The red line through the two images indicate the location of the A-scans. It seems as if the noise reduction has come at the prize of decreasing the height of the peaks. This is a consequence of the averaging of several noisy signals, but it is not know what the true peak heights are, or whether they are closer to the resulting peaks or not.

When looking at the sides of a macular hole in an average image, they appear slightly blurred, this is sometimes also the case with a cystic space. This is not due to a registration error, but because of small perpendicular shifts between the images in the set. These shifts often happen when dealing with macular holes, as previously explained, because visual acuity in the fovea is poor. It is also noticeable in one of the postoperative images, shown in figure B.5. This explanation corresponds well with the fact that this person has the worst visual acuity of the three postoperative cases tested.

Adjusting the parameters made no significant changes. This can be a sign of the robustness of the method, or that the task could be solved with a simpler method, or both. A less robust method of taking the maximum correlation for each A-scan would in some cases do satisfactory. Most of the columns would be registered correctly. One reason for this is that the used imagesets have fairly distinct layers. This is not the case for the imageset shown in Appendix C consisting of 11 images, where a less robust implementation fails in registering the images, but the method of tracking a shortest path in energyspace succeeds. So in the case of the macular hole images presented, it can be said that the task at hand is manageable and the method is robust, in a way that it can handle sets of images with less contrast than the ones presented. This is a quality that should not be underestimated, since other pathologies exists that leads to images with lower contrast than is the case for a macular hole.

Part III

Applications

Chapter 7

Pre-Operation

A macular hole typically develop through stages. The further a hole progress the worse the expected postoperative visual acuity is. But when for instance comparing two stage 3 holes, it is not well know what to look for. An obvious choice is the "size" of the hole. But what is a good measure of the "size"? Total volume, height, maximum width or something else? In [36] a few descriptors based on the configuration of the macular hole are investigated. The investigated descriptors were height, minimum width and base width of the hole. The measure that had the best correlation with postoperative visual acuity was the height of the hole divided by the base width of the hole.

There is a great interest in estimating these descriptors and others automatically or semi-automatically from an OCT image. If the retinal surface is found, the three sizes height, minimum and maximum width of the hole can be determined. The previously mentioned base diameter would very often be equivalent to the maximum width. But the difference between the two is that maximum width is well defined, compared to the base diameter, that needs to be measured at a small but self defined height above the base of the hole.

Another feature that can be calculated if the retinal surface is known, is the area of the hole, and thereby an estimate of the total volume as well.

Another interesting feature to determine is the thickness of the neuroretina. The neuroretina refers to the retina excluding the RPE. That means it starts at the transition between the vitreous body and the RNFL, and ends where the photoreceptors and the RPE connects. Since it can be very hard to locate the precise location where the RPE starts, I will measure the thickness down to the bottom of the outer nuclear layer, lying above the hyperreflective layer, ie. the IS/OS and RPE. In the following a reference to the neuroretinal thickness will refer to this thickness. Due to the difficulty in determining the true transition, this approximation is often used.

When having a macular hole, the neuroretina is often swelled with liquid around the fovea. The degree of swelling can be estimated from the neuroretinal thickness, but a value that gives an overall estimate is the neuroretinal area in the image.

When looking at an aligned OCT image of a normal individual, it has naturally been aligned to the RPE since it is the dominating reflecting layer. When looking at an aligned macular hole such as the one in figure 6.9 this is not the case. The alignment is a compromise between aligning the RPE and the outer retinal layers. If the image should be aligned to the RPE, it can be done if the RPE is located. The image should not be aligned to the top of the hyperreflective layer, since the IS/OS is present in the outer parts of the image, but not the foveal area. All the descriptors can be calculated and produced when the location of the top and bottom of the neuroretina are known. The central location of the RPE needs to be known as well, if an image aligned to the RPE should be produced.

7.1 Hyperreflective Layer

In this section an automated method to determine the center of the RPE and top of the hyperreflective layer will be described. It starts with finding the location of the center of the RPE with regularized dynamic programming. The top is found by searching in a modified gradient image with regularized dynamic programming above the central location.

Before any analysis is done, the image is gaussian filtered with a standard deviation of $\sigma = 1$. This reduces the noise, especially in the gradient image, which will be used later on. The location of the RPE can be found as the shortest path in the negated intensity image. There should be no chance of finding the location of the top of the retina, since that path would have a part going through background pixels where the hole is located, and generally the RPE has higher intensities than the top of the retina.

A typical OCT image, is aligned to the hyperreflective layer consisting of the RPE and IS/OS. This is because it is the dominant horizontal structure in the image, and several of the other layers lies parallel to it. This is not the case of images with macular holes. The alignment is a compromise between the RPE and the top layers of the retina, with the RNFL being the most reflective. This can be changed if the location of the image is aligned to the found central RPE location. The comparison between a macular hole image, and a normal or postoperative macular hole image should in this way be easier.

The top of the hyperreflective layer is found as the shortest path in the 50 pixels above the central line in the negated vertical gradient image. To allow several slopes the image has been reduced in its horizontal resolution, and the allowed vertical change has been increased, as described in section 4.1, before searching for the shortest path. The result is shown in figure 7.1, for one preoperative image and an average generated from 11 images, in the local area lying around the RPE. The top is marked with a red line. Most of the red line is not visible, since it is lying at the same location as a path that has been marked with black, which will be described shortly.

This procedure fails on some occasions. This is where it jumps down and locates the top of the RPE instead of the IS/OS. An extra energy term could be added that penalized lower lying pixels, but this could also shift the location of the shortest path upwards in the cases where the top has been correctly found. Therefore another approach is taken to fix this problem. The top of the IS/OS is expected to lie between 20 and 30 pixels above the central line, and the top of the RPE is expected to lie below 20 pixels. If there is a peak higher than 25% of the highest value in the gradient image, lying in the interval 20-30 pixels above the central line, it and its vertical vicinity is reduced in the space regularized dynamic programming searches, thus attracting the shortest path.

The path found searching in the negated gradient image where these rewards have been added, is shown with the black line. It can be seen in figure 7.1 that it pulls the path upwards to correctly find the top of the IS/OS on the right side in 7.1(a). But generally it is located at the same position as the original red line. If the method for some reason fails, guidance points that attracts the path are needed by the user.

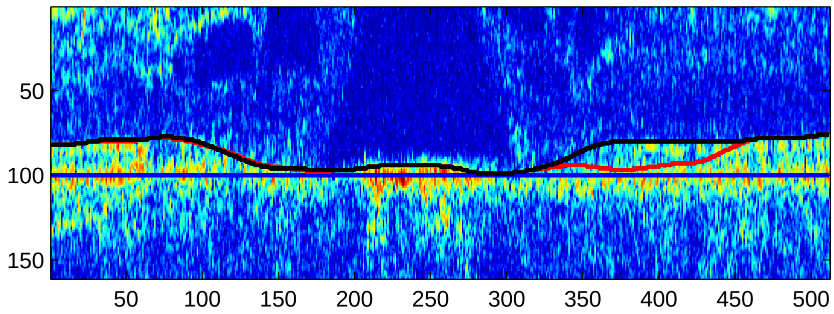

(a) Preoperative macular hole

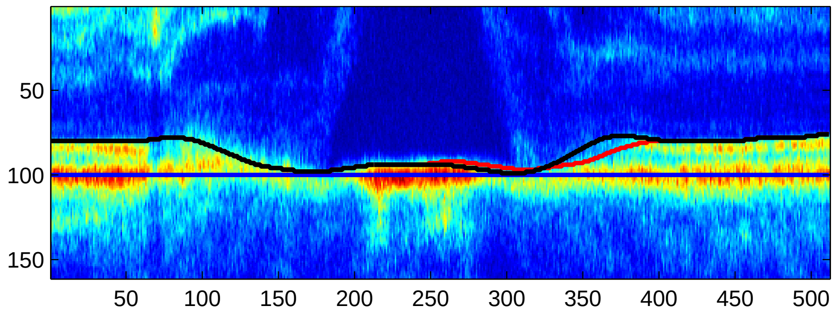

(b) Average of 11 images

Figure 7.1: Two sections of images with a macular hole. (a) is an original image, and (b) is one generated from 11 images. The central location of the RPE is marked with blue, and the shortest path found in the negated gradient image is marked with red. To correctly find the top of the IS/OS, pixels with a high gradient lying in the interval 20-30 pixels above the central line are rewarded.

7.2 Surface of Neuroretina

The path of a correctly found surface of the retina of a macular hole does not move from one column to the next as is the case when dealing with a retina with no macular hole. If the surface should be correctly found in the macular hole, the dynamic programming algorithm should not be used. A different approach is to use a parametric snake, or active contour.

An active contour is a curve that is initialized in the image, and deforms under the influence of forces. It is an iterative process, that falls into a local energy minimum. A traditional snake is a curve of length one, $v(s) = [x(s), y(s)], s \in [0, 1]$ that minimizes the sum of two types of energy

$$E_{snake} = \int_0^1 E_{int}(v(s)) + E_{ext}(v(s))ds. \tag{7.1}$$

The internal energy depends solely on the shape of the snake. It is often composed of a first and second-order term controlled by the two parameters α and β that does not have to be constants, but is often chosen as such, and an extra constraint term. The only extra constraint term I will mention is a balloon force that for a closed curve exerts a force in the normal direction of the curve. This could for instance ensure that the curve does not shrink into a point, or ensure a higher captive range. The force is controlled by the constant γ.

$$E_{int} = \frac{1}{2}(\alpha|v'(s)|^2 + \beta|v''(s)|^2) + \gamma n \cdot v(s) \tag{7.2}$$

The external energy can depends on the image values. It is often constructed such that the snake is attracted to edges, where the numerical gradient or squared gradient are obvious choices. The gradient is often calculated on a gaussian filtered image to reduce the level of noise in the gradient, and improve the snakes captive range. The force is controlled by the constant δ.

$$E_{ext} = -\delta|G_\sigma * \nabla I(x, y)|^2 \tag{7.3}$$

For further information and how to implement a snake see for instance [37].

Since the surface of the retina is not a closed curve, the snake is adapted, such that it always starts at one side of the image and ends at the other side. This way the snake will never shrink into a point, but the balloon force is included anyway to ensure that the snake extends to all the corners in the macular hole. This is done by letting the balloon force be directed perpendicular to the snake in the downwards direction. This only gives good results if the snake is initialized above the border it should fall to rest on.

In figure 7.2 two different initializations and the resulting snakes are shown. With the same choice of parameters, the snake seem to fall to rest at the same location. It is also possible to generate a snake that falls to rest correctly in the macular hole from a straight line lying above the retina, but this requires a high γ-value, which makes the final location of the snake in the macular hole extends into the retina on the sides fairly often. An initialization as in 7.2(a) or better is recommended.

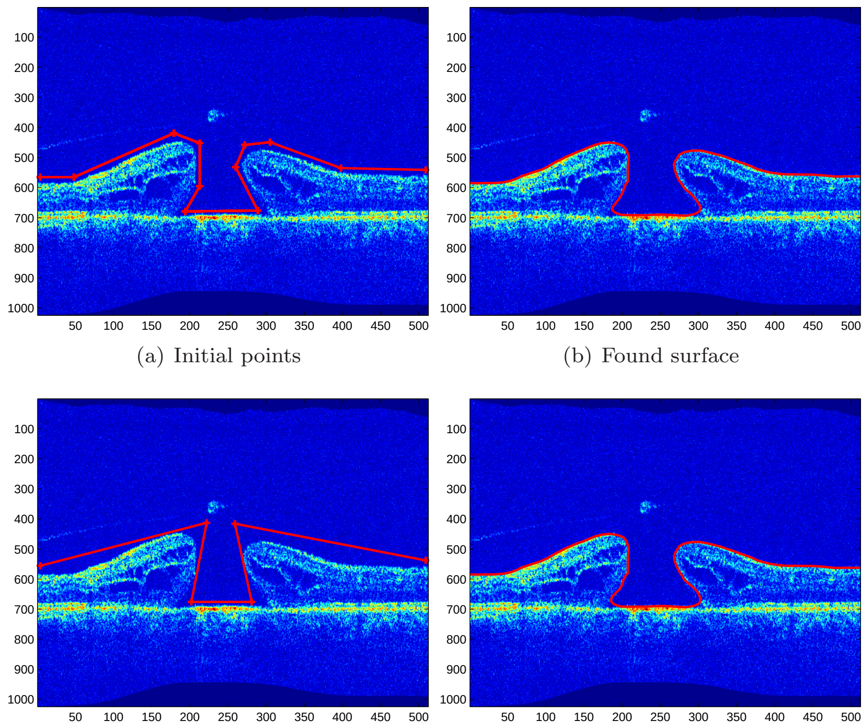

(a) Initial points (b) Found surface

(c) Initial points (d) Found surface

Figure 7.2: Two different initializations are shown to the left, and the resulting snakes are shown to the right. Both initializations gives a satisfactory result. If the parameters are altered significantly, the edges of the hole are not found correctly though.

A classic problem with snakes is that there are so many parameters to adjust. this is also the case with this implementation, where five parameters exists besides the initialization. But once the method has been adjusted once to an OCT image, limited adjustment is needed.

The two cases from figure 7.1 are shown with the neuroretina marked and aligned to the center of the RPE in figure 7.3. The average image is the image used for determining the hole descriptors the in next section.

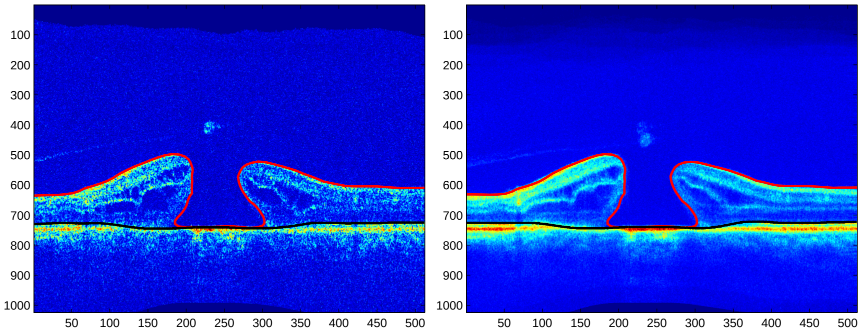

(a) One image of macular hole (b) Average image of macular hole

Figure 7.3: The images have been aligned to the central part of the RPE instead of the standard procedure with maximizing correlation. This way comparison with a non-macular hole image is easier, since the alignment is the same. The top and bottom of the neuroretina have been marked in the images.

7.3 Determining Hole Descriptors

Once the location of the top and bottom of the neuroretina have been found it is just a matter of writing a robust code that can handle the variation in the shapes to extract the hole descriptors. I will describe the principles I have used to find the descriptors. In the following will be referred to figure 7.4, where the descriptors are shown.

The first thing determined is the point of overflow in the hole. This is marked with the green horizontal line. This is the lowest of the highest points on the left and right side. Once this point has been determined, the width of the hole as a function of the row can be calculated. The minimum is marked with the top horizontal arrow, and the maximum that must lie below the minimum, is marked with the bottom arrow. The height of the hole can vary depending on where it is measured, so it is decided that it is measured at the central point at the minimal width of the hole. This way a point to center the hole has also been determined.

In this case the overflow happens at row 522, the minimum and maximum width are 60pixels = 0.70mm and 117pixels = 1.37mm, the height of the hole at the central row 238 is 217pixels = 0.42mm.

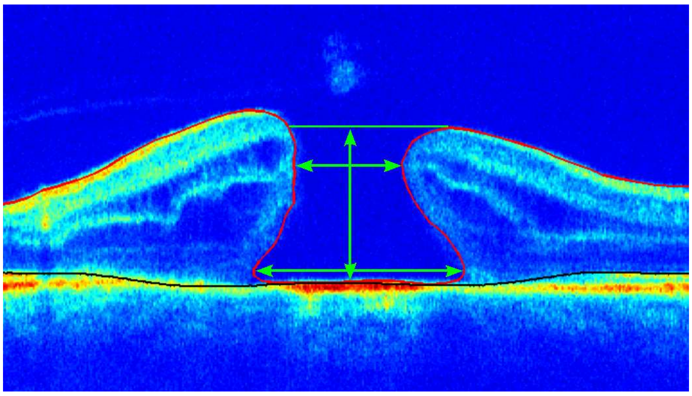

Figure 7.4: The point of overflow is marked with a horizontal line. The three hole descriptors marked with arrows, are the minimum width, the maximum width and the height of the hole.

From the surface of the retina and the overflow point, the area inside the macular hole can be determined as the points lying inside the polygon consisting of the points on the surface starting and ending at the vertical position of the point of overflow. The area has been marked with black in figure 7.5. In this case the area of the macular hole is 0.41mm^2. The neuroretinal area is 1.54mm^2. This is not an extreme value for the neuroretinal area, but considering the macular hole is present, it means that the neuroretina outside the macular area is significantly swelled.

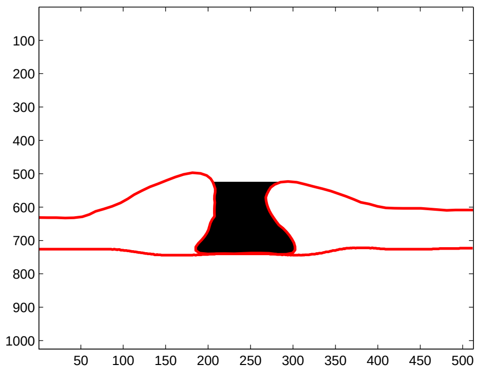

Figure 7.5: The top and bottom of the neuroretina is outlined, and the area of the macular hole has been marked with black. In this case the area inside the hole is 0.41mm^2.

From the area of the hole, an estimated volume can be calculated. It is calculated as if the hole is circular symmetrical about the central point, and each side contribute with π radians.

With cylindrical coordinates the estimated volume can be calculated as

$$\begin{aligned} \mathrm{Vol}_1 &= \int_A 1 d\Omega = \int_0^{\frac{\pi}{2}} \int_0^h \int_0^{r_1(z)} \varrho \, d\varrho dz d\phi + \int_{\frac{\pi}{2}}^{\pi} \int_0^h \int_0^{r_2(z)} \varrho \, d\varrho dz d\phi \\ &= \frac{\pi}{2} \int_0^h r_1(z)^2 + r_2(z)^2 dz \end{aligned} \tag{7.4}$$

where h is the height of the macular hole and $r_1(z)$ and $r_2(z)$ are the distances from the center to the left and right side respectively. This gives a total estimated volume of the hole of $1.23 \cdot 10^6 \text{voxels} = 0.33\text{mm}^3$

This can be extended to give a better estimate of the volume, if the shape and size of the hole is known for six radial scans through the fovea at different angles, as will be shown in section 9.2.

7.4 Thickness of Neuroretina

For a pathology called macular edema, where a swelling is occuring inside the retina, a good way to get an overview of the extent of the edema is to plot a 2D surface that represents the thickness of the neuroretina. This can be generated from a set of scans into the retina, where the thickness of the neuroretina has been outlined on each. A typical procedure is six radial scans running through the fovea. Looking at a set of retinal thickness', it is not easy to visualize the actual location and extent of the edema, a surface plot helps with this.

Since a macular hole often is accompanied with a significant swelling in the retina, such a visualization may be of interest. It has been decided that the neuroretinal thickness should be represented as zero where it has been detached from the RPE. This way there is no ambiguity of where the detachment is happening. The retinal thickness for the average image of the macular hole is plotted in red in figure 7.6. The original shape of the retina can be seen in green. It lies behind the red curve outside the hole. A normal retinal thickness is plotted in blue. The thickness is measured in pixels. The swelling in the macular area is so significant, that the neuroretina has almost doubled its thickness at the thickest points.

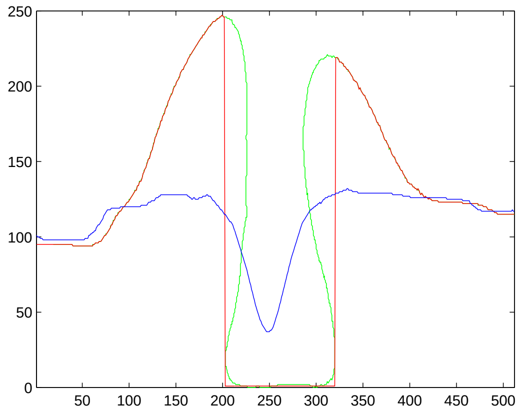

Figure 7.6: The retinal thickness for the average image of the macular hole is plotted in red. The original shape of the retina can be seen in green. It lies behind the red curve outside the hole. A normal retinal thickness is plotted in blue. The thickness is measured in pixels. The swelling in the macular area is so significant, that the neuroretina has almost doubled its thickness at the thickest points.

If the neuroretinal thickness is found for a set of images running through the fovea, a 2D-surface of the thickness can be generated as will be investigated in section 9.1.

To sum up, it is possible to automatically determine the top of the hyperreflective layer and the top of the neuroretina semi-automatically. From these all relevant hole descriptors can be determined. The procedure can be done on a single image, but using an averaged image instead makes the method more robust.

7.5 Pre-OMaH Tool

The procedure to extract the relevant preoperative descriptors have been implemented in Borland C++ Builder. A screen shot is shown in figure 7.7. The top of the hyperreflective layer has been found automatically, and the top of the neuroretina has been found on the basis of ten points provided by the user. Once these two lines have been found, the program determines the height and the minimum and maximum width of the hole, the area of the retina, the area of the hole, and the estimated volume of the hole.

The two main users opinion about the two programs developed can be read in section 8.4, where the postoperative program is presented.

Some of the variables have been locked at appropriate values for clearness. The only variable that can be altered by the user are β, γ and δ for the active contour when locating the top of the retina, and a smoothing parameter for the top of the hyperreflective layer. A short user manual for Pre-OMaH Tool is provided in Appendix F, and a copy of the program is provided on the enclosed CD along with two image examples.

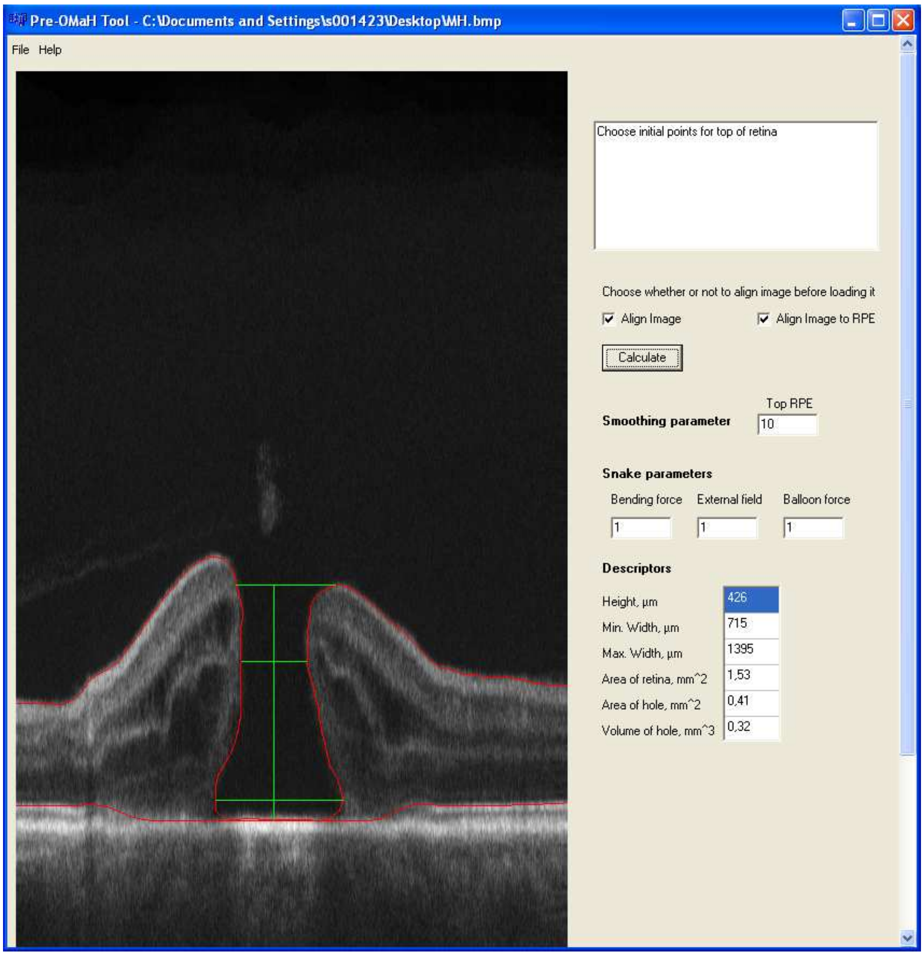

Figure 7.7: Screen shot from Pre-OMaH Tool. All the relevant descriptors have been determined.

Chapter 8

Post-Operation

If a macular hole is operated, the custom procedure used at Herlev Hospital is to have a follow up inspection of the eye after 3 month and then again after 6 and 12 month. At each of these consultations the visual acuity is measured, and a couple of examinations are performed, among these is the taking of OCT images. Two procedures are recorded. One is 6 radial scans through fovea, with $\frac{\pi}{6}$ radians in between. This gives a good overall impression of the foveal area. And the other is taking several images at the same location through the fovea. This is for producing one combined image with reduced noise.

One follow up OCT image is shown in figure 8.1(b), where comparison to a normal subject, shown in figure 8.1(a), is possible. There are several things to look for when assessing the outcome of the operation. First of all it should be mentioned that the hyperreflective top layer on the right of the normal individual is not one. This is the RNFL, and it is thicker here because it is close to the optical nerve head. That means the image is taken at an angle close to horizontal, and the thickened RNFL indicate the nasal direction.

Three factors that may indicate a loss of visual acuity in the fovea will be mentioned in the following. These factors can be seen even more easily on the averaged image consisting of the image in figure 8.1(b) and 13 others taken of the same retinal area shown in 6.10(b).

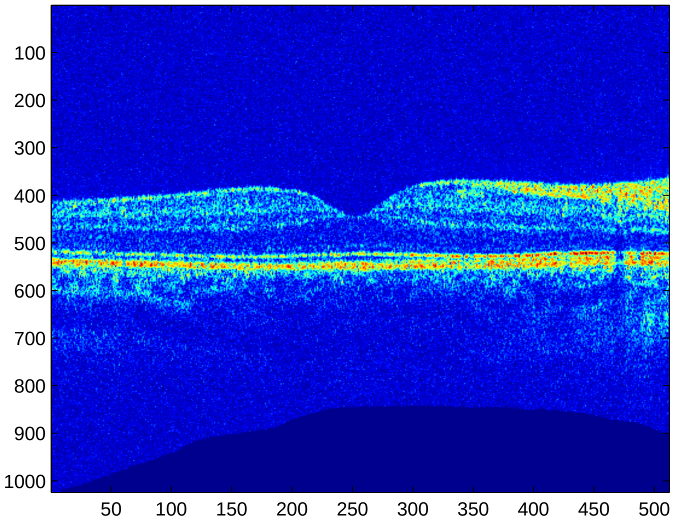

(a) Normal eye

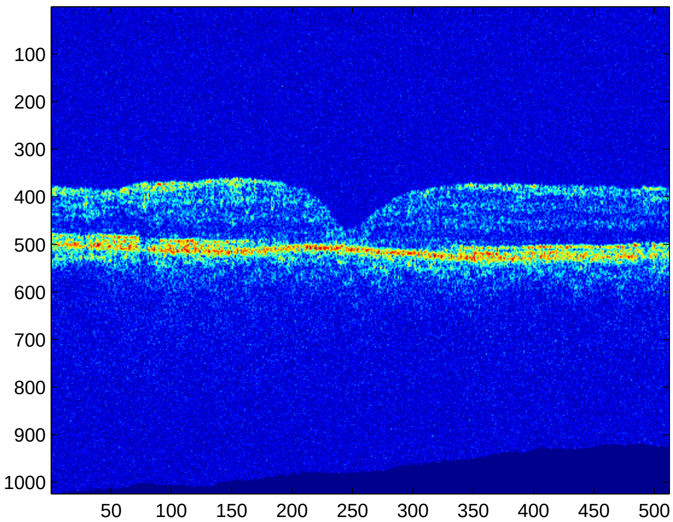

(b) Postoperative eye

Figure 8.1: The images shows an OCT image of a normal individual and one having had macular hole surgery. Three factors that indicate loss of visual acuity in the fovea are: Thinning of the hyperreflective layer, larger foveal dip and increased intensity of the outer nuclear layer in the fovea.

It can be seen that the hyperreflective layer in the fovea has decreased in thickness. It consists not only of the RPE as previously mentioned, but also of the IS/OS. The IS/OS is the thin hyperreflective layer at the top that in the normal eye separates itself from the RPE in the foveal area. In the postoperative image this layer has completely disappeared in the foveal area. This indicates that the photoreceptors may have died in the detachment period.

The "foveal dip" is significantly deeper in the postoperative image. Since the significant layer present in the fovea lying above the IS/OS is a part of the photoreceptors called the outer nuclear layer, it indicates that not only the IS/OS has disappeared but the photoreceptors have also shrunk or completely disappeared. This of course is not a good sign. These two features have been investigated before in [38], where there were indications of the IS/OS thickness correlated with visual acuity but not the foveal thickness.

The last factor that indicates a significant loss of visual acuity in the postoperative image is the increased intensity of the outer nuclear layer in the fovea. It can be seen, and even better in figure 6.10(b) that the layer in the fovea lying above the RPE has a heightened intensity when compared to the non foveal area in the same image, or the foveal area in the normal eye. This could mean that instead of being rejuvenated photoreceptors it may be scar tissue.

It would be beneficial to be able to determine these three factors quantitatively, preferably automatic. The next sections describes a procedure that does just that.

8.1 Thickness of Hyperreflective Layer

In this section an automated method to determine the top and bottom of the hyperreflective layer will be described. It is very similar to the method in section 7.1, but a few differences exists. An approximate location of the hyperreflective layer can be easily found as the maximum of the horizontally summed image. A more accurate location is determined with regularized dynamic programming. The actual top and bottom is found by searching in the gradient image with regularized dynamic programming above and below the central location.

In figure 8.2 is shown the horizontal sum of the two images in 8.1. The blue is the normal eye. It can be seen the first peak is wider for the normal individual due to the RNFL. But the height of this peak does not even get close to the height of the peak corresponding to the RPE. An enhanced RNFL will of course never be present in both sides of the image, so a robust initial guess of the location of the center of the RPE in an aligned OCT image of the retina is at the maximum of the horizontally summed intensities. This could not be done on the macular hole images, since they were not aligned with the hyperreflective layer.

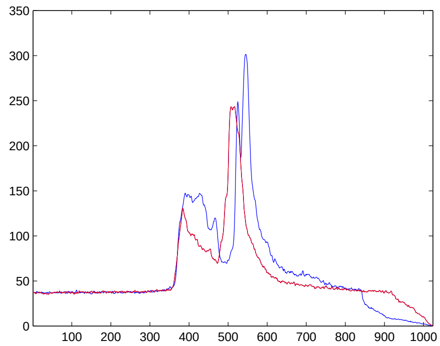

Figure 8.2: The two graphs are the horizontal sums of the pixel values of the two images in 8.1. The blue is the normal eye. A robust initial guess of the location of the center of the RPE is at the maximum of the horizontally summed intensities.

Before further analysis is done, the image is gaussian filtered with a standard deviation $\sigma = 1$. The horizontal estimate is used to locate a part of the image that contains the hyperreflective layer. In this part of the image, regularized dynamic programming is applied to find the shortest path in the negated images, such that the solution found is the approximate location of the center of the RPE. The reason for not using the horizontal line to divide the top from the bottom is that the image may not be perfectly aligned to the RPE and the reason for using it at all is that it speeds up the process, when an approximate location of the RPE is known.

The central line found this way is shown for three cases in 8.3. The line is shown in blue, and the horizontal line found as the maximum of the horizontal sum of the image is shown in red.

Once the central line is found the image is split in two parts one above and one below the line. The top of the hyperreflective layer is found as the shortest path in the negated vertical gradient image above the line, and the bottom is found as the shortest path in the vertical gradient image below the central line. The result is shown in figure 8.3 marked with black lines. The result can be seen on the full image on figure 8.5 as well. To allow several slopes the image has been reduced in its horizontal resolution, as described in section 4.1.

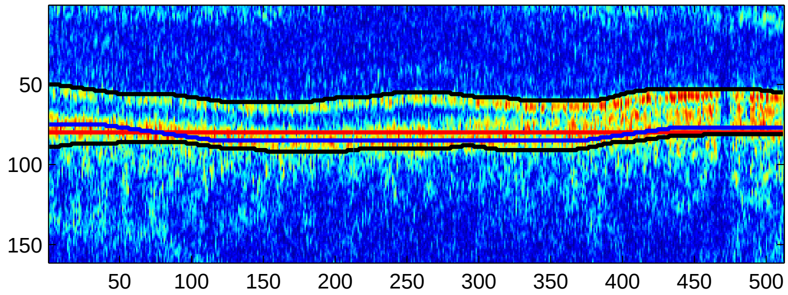

(a) Normal eye

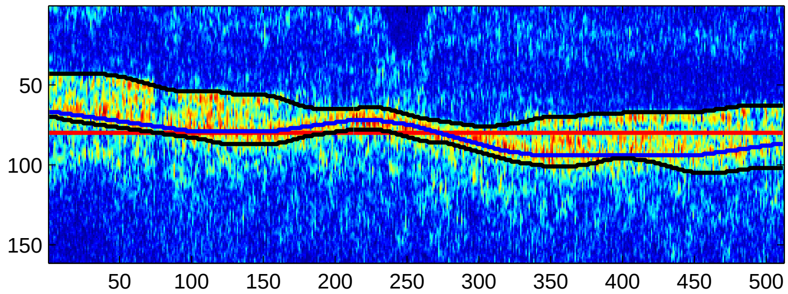

(b) Postoperative macular hole

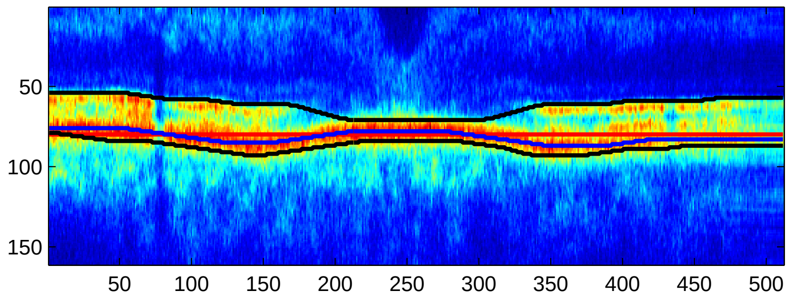

(c) Average of 14 postoperative macular hole

Figure 8.3: Three cases where the thickness of the hyperreflective layer should be located. The initial horizontal line found as the maximum of the horizontal sum of the image is shown in red. The central line located in the area around the red line is shown in blue. The top of the hyperreflective layer is found as the shortest path in the negated vertical gradient image above the line, and the bottom is found as the shortest path in the vertical gradient image below the central line. They are both shown with black lines.

Looking at 8.3(b), it is obvious that the initial horizontal line can not be used for dividing the image in two and expect to have the bottom edge of the RPE lying below, since the black line crosses the red line twice. This happens at the far left side for the averaged image as well. This image has been better aligned, because of the less noise present in the image. There are places, where the bottom line gets very close to the blue line as well, but if the shape constraints are not considered, it should never reach it, since the blue line indicates the pixel maximum, and the black the minimum of the vertical gradient.

In figure 8.4, the three corresponding thickness' of the hyperreflective layer are plotted. The thickness measured on the normal eye is plotted in green, from one image of the postoperative macular hole is plotted in blue and from the averaged postoperative macular hole is plotted in red. As expected the red and blue has a dip in the center. They follow each other very well, except for the right side, where the blue increases significantly. This can be observed on the image as well, where the bottom line lies too far down. This is not a significant error for two reasons. First of all the region of interest is the center compared with the thickness lying up to 2μm (170 horizontal pixels) from the center and the method is primarily expected to be used on averaged images, where it has no problem locating the edges.

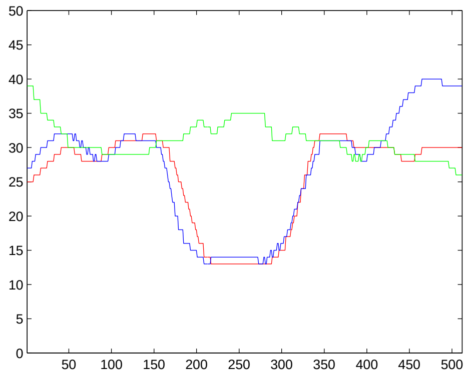

Figure 8.4: The thickness of the hyperreflective layer found for the three cases in figure 8.3 are plotted in pixels. The thickness measured on the normal eye is plotted in green, from one image of the postoperative macular hole is plotted in blue and from the averaged postoperative macular hole is plotted in red. The red and blue follow each other very well, except for the right side, where the blue increases significantly.

The only thing it can affect is when trying to estimate the width of a possible atrophy, ie. shrinkage of the hyperreflective layer. A simple way is to calculate the average thickness of the outer one mm, and define a location as atrophic if it is less than two thirds of this value. This gives an estimate width of the atrophy as 1.41mm for the average image and 1.75mm from one image, where the discrepancy is of course due to the difference in the right side.

A problem that could be expected to occur in some images is that the path may lie at the top of the RPE instead of the IS/OS, as happened with the macular hole. The same approach as was used in that case could be implemented, but the necessity of this has so far seemed to be minimal.

8.2 Thickness of Neuroretina

Since the transition between the photoreceptor layer and the IS/OS already has been determined, all that needs to be done to find the thickness of the neuroretina is to locate the top of the neuroretina. This can be achieved the same way as finding the top and bottom of the hyperreflective layer, by use of regularized dynamic programming. When locating the top of the neuroretina, a larger vertical change must be allowed, to properly locate the foveal dip. The result of searching the negated gradient of the image above the top of the hyperreflective area with an allowed vertical change of three, and subsampling horizontally by a factor of two is shown in figure 8.5. The previously determined borders lying at the top and bottom of the hyperreflective layer are also shown. The result is very satisfactory even on an image that is not an average of several images.

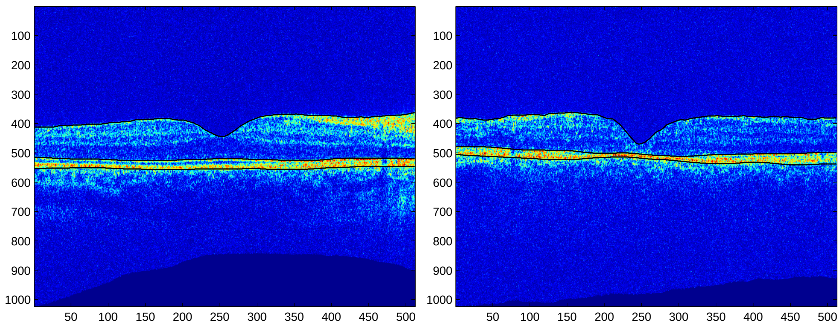

(a) Normal eye (b) Postoperative macular hole

Figure 8.5: The top of the neuroretina is outlined, along with the top and bottom of the hyperreflective layer. The hyperreflective layer consists of the IS/OS and the RPE.

The thickness of the neuroretina for the same three cases used in figure 8.4 are plotted in 8.6. The normal case in green has a significantly thicker foveal dip, as was visually observed as well. The thickness of the neuroretina from the averaged image and that of one image lies very close as hoped. The thickness measured from one image is less smooth because of the significantly higher noise level in the image. But it seems as if it is not necessary to average several images to measure the thickness of the neuroretina.

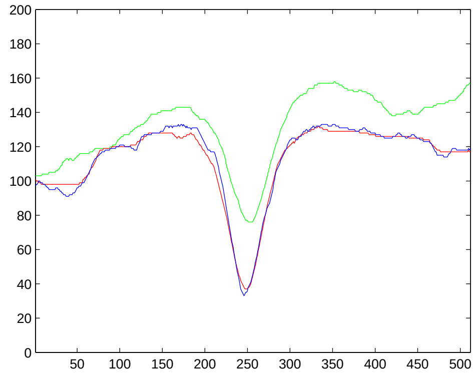

Figure 8.6: The thickness of the neuroretina for the same three cases used in figure 8.4. The normal case in green has a significantly thicker foveal dip, as was visually seen as well. The thickness of the neuroretina from the averaged image and that of one image lies very close as hoped. The thickness measured from one image is less smooth because of the significantly higher noise level in the image.

Another descriptor that can be evaluated is the neuroretinal area. It indicates whether the neuroretina has swelled or shrunk. The estimated areas are for the normal case 1.54mm^2, the average image 1.32mm^2 and the one image 1.33mm^2. The postoperative retina has shrunk significantly as can also be observed in figure 8.6.

If the neuroretinal thickness is found for a set of images running through the fovea, a 2D-surface of the thickness can be generated as will be investigated in section 9.1.

8.3 Intensity of Outer Nuclear Layer

When the top and the bottom location of the neuroretina is known, it is fairly straightforward to sample the outer nuclear layer. The only thing that needs to be decided is how wide the sampling should be, and how close to the hyper-reflective layer the sampling should start. It should not start right above the transition line, first of all to get away from the gradient located at the edge, but also because a thin layer called the external limiting membrane lies between the IS/OS and the outer nuclear layer. It is only slightly more reflective than the outer nuclear layer, but never the less it should not be sampled.

It is estimated that the external limiting membrane is approximately 10 pixels wide, so starting the sampling 15 pixels above the transition line is appropriate. A sampling height of 15 pixels gives an acceptable compromise between not entering the next layer (the outer plexiform layer), and sampling a sufficient amount of pixels to reduce noise. If the foveal dip is so deep that it enters the sampled area, samples are only taken for those pixels belonging to the neuroretina of course. The average intensity for each 15 vertical pixels selected as a function of column number is shown in figure 8.7(a) for the averaged and non-averaged image of a postoperative macular hole. The normal case is not plotted for clarity, but it is approximately as noisy as the blue line and fluctuating around 0.15. The non-averaged image seems to be a noisy version of the averaged image.

In figure 8.7(b) the intensities are shown, after filtering with a gaussian filter with $\sigma = 10$. The non-averaged approaches the averaged, but there is still a higher degree of fluctuation. A common characteristic for the two curves are that they peak where the foveal dip is located in the image.

To sum up, it is possible to automatically determine the thickness of the hyperreflective layer and the neuroretina as well as the local intensity of the outer nuclear layer for a normal retina or a postoperative macular hole. It can be done on a single image, but using an averaged image instead reduces the noise in the results significantly.

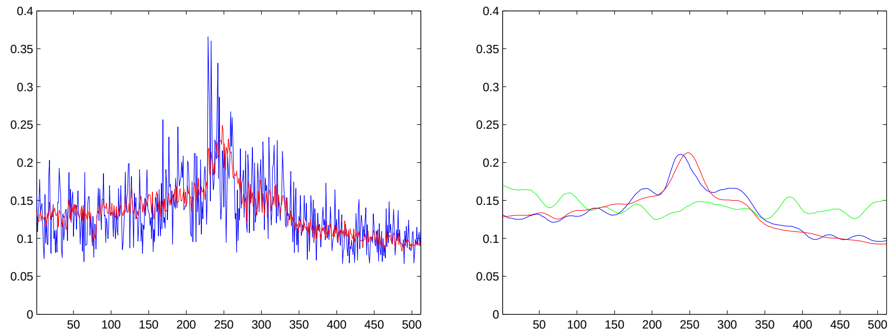

(a) Intensities (b) Intensities filtered

Figure 8.7: The intensities of the outer nuclear layer for the same three cases used in figure 8.4. (a) shows the measured intensities where each point on the graph is an average of 15 pixels. The normal case is not plotted for clarity. The non-averaged image semm to be a noisy version of the averaged image. In (b) the intensities are shown, after filtering with a gaussian filter with $\sigma = 10$. The non-averaged approaches the averaged, but there is still a higher degree of fluctuation. A common characteristic for the two curves are that they peak where the foveal dip is located in the image.

8.4 Post-OMaH Tool

The procedure to extract the relevant postoperative descriptors have been implemented in Borland C++ Builder. A screen shot is shown in figure 8.8. All the three layer have been found automatically. On the basis of these three lines, the thickness of the neuroretina and hyperreflective layer are sampled at predefined locations, and the width of a possible atrophy is estimated along with the neuroretinal area. The intensity of the outer nuclear layer is estimated in the fovea and on the side by sampling two squares instead of producing a graph. The two squares are shown in the image with a slight green and red enhancement.

The only variable that can be altered by the user are smoothing parameters for each of the three lines. But the main way to affect the location of transitional layers are off course by clicking on the image, and thereby lowering the path cost at and around the given location. The program has been tested on 6 different images, from 6 different postoperative cases. The results can be seen in Appendix E. A short user manual for Post-OMaH Tool is provided in Appendix G, and a copy of the program is provided on the enclosed CD along with two image examples.

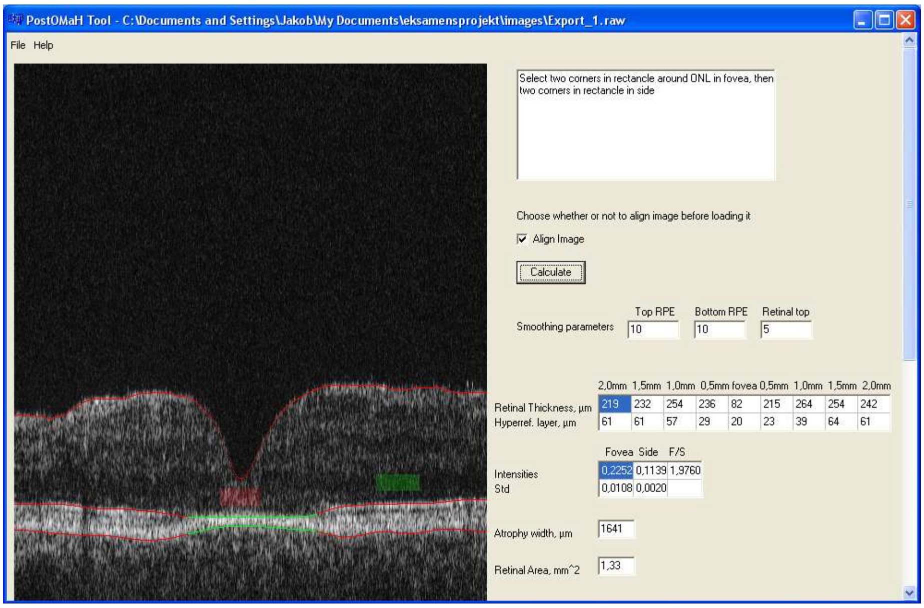

Figure 8.8: Screen shot from Post-OMaH Tool. All the relevant descriptors have been determined.

The two main users of Pre-OMaH and Post-OMaH Tool have been asked for their opinion about the two programs. The comment received about the applicability and prospects of the programs was:

”Programmerne har gjort det muligt at hente langt flere oplysninger ud fra helt almindelige undersøgelsesdata end det hidtil har været muligt med det kommercielt tilgængelige software. Således har programerne gjort det muligt at kvantificere en række forhold i nethinden som ikke tidligere har været gjort til genstand for systematisk undersøgelse[1].”

The comment received about the actual implementation was:

”De omfattende algoritmer er gjort tilgængelige i et brugervenligt program som efter kort introduktion kunne anvendes[2].”

[1]In English: The programs have made it possible to extract more information from typical examination data than so far have been possible with commercially available software. Thus, the programs have made it possible to quantify a number of conditions in the retina that previously have not been the subject of systematic examination

[2]In English: The extensive algorithms are made available in a userfriendly program that after a brief introduction could be used

Only positive response was received, so it seems as the dialogue throughout the project from when deciding which descriptors should be determined to the introduction to the actual implementations has been satisfactory.

Chapter 9

Beyond One Slice

Up to this point the actual location of the scans investigated have not been used. This is because if just one scan is available it is as good as impossible to say anything about the part of the retina lying outside the scan. If several scans are available the relative location between them becomes of interest, so features of the intermediate tissue can be estimated by interpolating features from the scans.

This approach will in the following sections be applied to six radial scans running through the fovea. This is a standard scan procedure in the StratusOCT apparatus, but the general ideas can of course be applied to any number of scans of any shape.

9.1 Surface Mapping

It can be very useful to generate a surface map representing the neuroretinal thickness for a large part of the back of the eye in the case of a pathology called macular edema. When looking at the retina through the lens it can be hard to see which areas are swelled. With a surface map of the neuroretinal thickness the ophthalmologist can visualize the size, shape and location of the swelling. When focusing on applications within macular holes, this can be used to visualize the swelling around a macular hole, but it can also be used to evaluate if the neuroretinal thickness postoperative has dropped below normal in the foveal area.

The method to generate the surface will not be discussed in detail, it will just be briefly explained in the following. First the neuroretinal thickness is found for a set of radial scans running through the fovea, normally six, with a difference in angle of $\frac{\pi}{6}$ radians. These points are mapped back to their original position on the retina. So a surface needs to be generated from a set of known thickness points located as shown in figure 9.1.

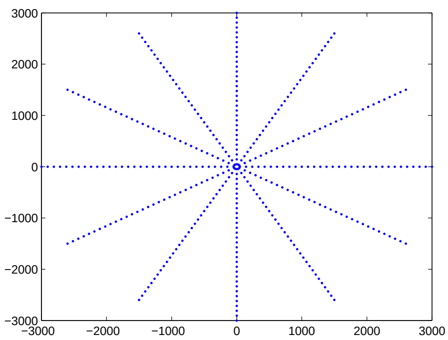

Figure 9.1: Known point locations projected on a plane.

This can be done by triangulating the convex hull, ie. the smallest convex region enclosing all the points. A standard method is the Delaunay triangulation. The points lying within a given triangle are then interpolated from the known corner points. This method is not very robust to noise though, since the surface is forced through all points. As will be seen shortly this is generally not a problem in our case. But if the scans are not aligned at the intersections of the scan lines, the method produces almost useless results.

The surface of the neuroretinal thickness is shown for four different cases in figure 9.2. The colormap used is the same as used by Zeiss in their StratusOCT. Since it is so commonly used by ophthalmologists it has become their colormap of choice. It is shown below the surface maps. The diameter for all surfaces is 6mm. Figure 9.2(a) shows a normal eye with the fovea in the center, and a slightly thicker right side than left, indicating that this is the side the optical nerve head is located in the surface map. Figure 9.2(b) shows a macular edema, where it can be seen that there actually are two different swellings. One is getting very close to the fovea, and the other is slightly smaller and further away.

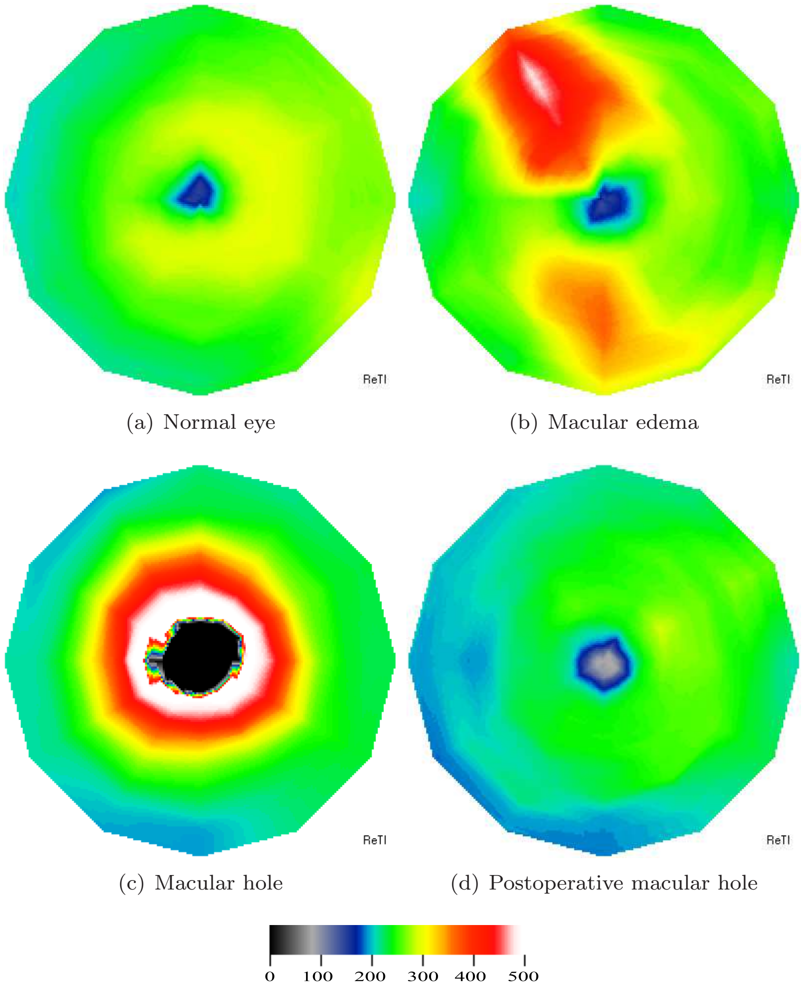

(a) Normal eye (b) Macular edema

(c) Macular hole (d) Postoperative macular hole

(e) Colormap in μm

Figure 9.2: The surface of the neuroretinal thickness is shown for four different cases. (a) shows a normal eye with the fovea in the center. (b) shows a macular edema, where two different swellings are apparent. A macular hole is shown in (c), where the thickness has been set to zero for places where the retina is detached from the RPE. There is a significant swelling outside the detached area. The last case in (d) is a postoperative image of a macular hole. It seems as if the entire area is thinner than the normal case, especially the fovea.

A macular hole is shown in 9.2(c), where the thickness has been set to zero for places where the retina is detached from the RPE. This creates artifacts at the borders. The significant swelling around the hole is apparent, but the only information you can get about the actual hole from this is the size.

The last case shown in 9.2(d) is a postoperative image of a macular hole. The fovea is significantly thinner, and it seems as if the entire area is thinner than the normal case. Whether this is significant of course needs further investigation, but the overall impression is a depleted neuroretina.

If the top border of the retina and the transition between the the outer nuclear layer and the IS/OS are known, the thickness of the neuroretina can be plotted. If it has been found for a set of radial images, a surface can be interpolated. It is very useful in the case of macular edemas, where one case has been shown, but it does not provide any information about a macular hole except the size of the detachment. If the surface of the neuroretina is reconstructed in 3D, more information would be available as will be shown in the next section. Postoperative cases can be investigated for whether the foveal pit lies too deep in the retina, and if the entire neuroretina is depleted.

9.2 3D Reconstruction

Another way to visualize the neuroretinal thickness in figure 9.2 is to plot a surface in 3D. In the normal or postoperative case, it is a different way of representing the same data, but in the case of a macular hole, where the thickness no longer is a function of the location on the retina, it gives a huge advantage, since the shape of the hole can now be visualized.

To generate the surfaces, a method called radial basis function has been used, it is a natural way to interpolate scattered data, particularly when the data samples do not lie on a regular grid and when the sampling density varies [39].

A Toolbox for Matlab from FarField Technology called FastRBF has been used to generate the surfaces. Three cases can be seen in figure 9.3. The same colormap used in figure 9.2 has been used to represent the neuroretinal thickness. The points used to generate the surfaces has been plotted with blue crosses. When comparing the normal eye to the postoperative macular hole, the same observations as was made when looking at the surface map can be seen. The foveal pit is deeper for the postoperative case, and the entire surface seem depleted.

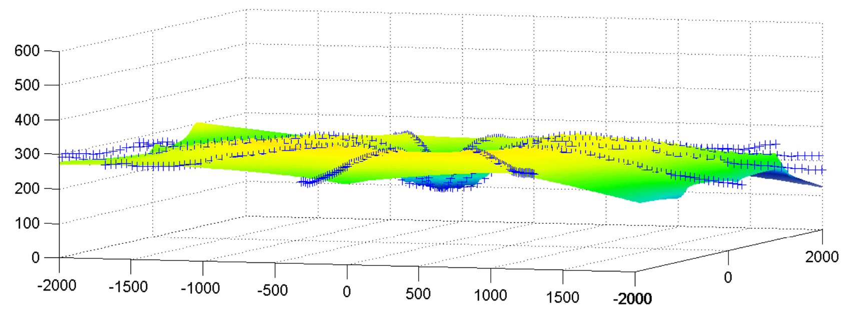

(a) Normal eye

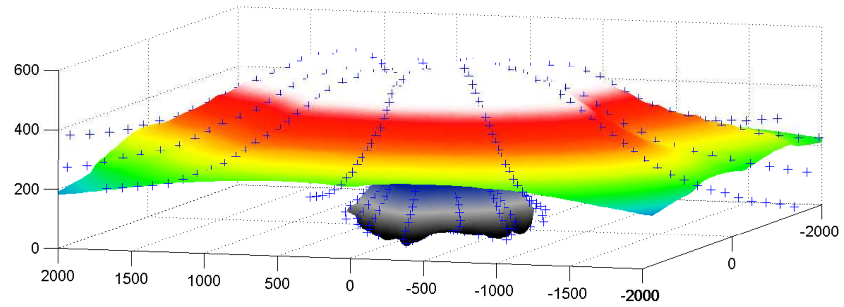

(b) Macular hole

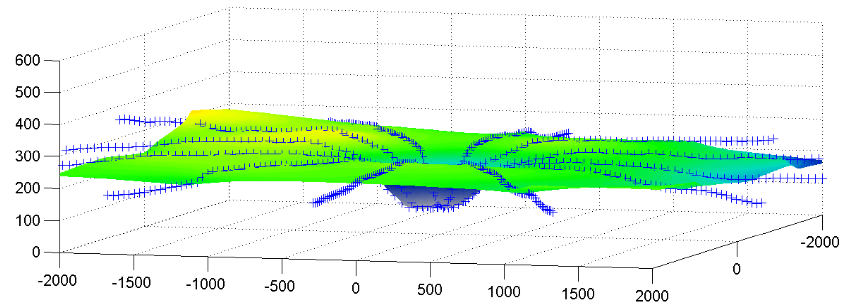

(c) Postoperative macular hole

Figure 9.3: The points used to generate the surfaces has been plotted with blue crosses. When comparing the normal eye to the postoperative macular hole, the same observations as was made when looking at the surface map can be seen. The foveal pit is deeper for the postoperative case, and the entire surface seem depleted. When looking at the macular hole, the bottom shape of the hole can be observed.

When looking at the macular hole, the bottom shape of the hole can be seen in 9.3(b). If the point of view is moved downward, the entire shape of the hole can be seen as shown in figure 9.4. The minimum and maximum size of the hole can be observed when inspecting the surface from top and bottom respectively, but the 3D cup is now also available, which has not been the case in any of the previous methods.

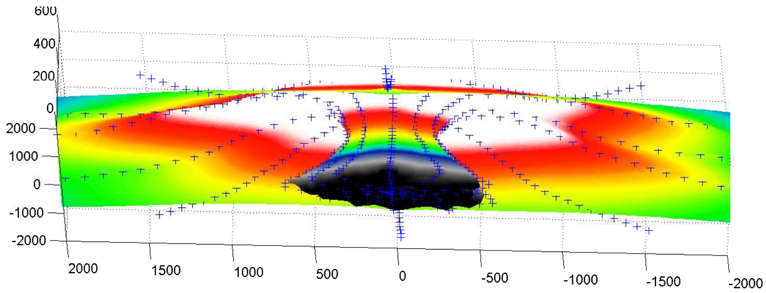

Figure 9.4: The point of view has been moved downward, when compared to 9.3(b), so the entire shape of the hole can now be seen.

In equation 7.4 the volume of the hole was estimated from one scan. Since six radial scans and six corresponding cup shapes now are available, a better estimate of the volume can be made. If the same approach is used, by assuming circular symmetry about the central point, and letting each side contribute with $\frac{\pi}{6}$ radians, the estimated volume is going to be

$$\mathrm{Vol}_6 = \frac{\pi}{12}\int_0^h \sum_{i=1}^{12} r_i(z)^2 dz = \frac{1}{6}\sum_{j=1}^{6} \mathrm{Vol}_1(j) \tag{9.1}$$

Where j is the radial scan number. So the estimate based on 6 scans is just the average of the six volume estimates based on one scan at a time. For the case visualized $\mathrm{Vol}_6 = 0.29\mathrm{mm}^3$ with a standard deviation of 0.07 among the six Vol_1 estimates. This is a fairly high standard deviation, which can be attributed to two types of variation. The first being actual shape variation in the macular hole, and the other being due to off-centered scans and not correctly found surfaces. An off-centered scan would lead to a larger estimated volume. The volume estimate could therefore be expected to be slightly biased, if the scans are not correctly centered.

The neuroretinal thickness can be visualized in 3D by a surface. In the normal or postoperative case, it is a different way of representing the same data, but in the case of a macular hole it gives a huge advantage, since the shape of the hole can now be inspected. The six available scans also gives a better estimate of the volume of the cup.

Part IV

Discussion

CHAPTER 10

Discussion

In the next sections part II and III of the project will be summarized and concluding remarks are given.

10.1 Image Enhancement

Imaging methods relying on measuring coherent signals can be affected by interference and therefore speckle. In chapter 5 three different types of diffusion were applied to OCT images, to examine how effective they each were to reduce speckle.

The simplest type was linear isotropic diffusion, which is equivalent to a convolution with a Gaussian kernel. The result is of course a gradual smoothing of the initial image. The effect of speckle disappears, but so do details in the image. Very similar results were seen with complex diffusion. With coherence enhancing diffusion, smoothing was sometimes done perpendicular to the actual layers, which lead to new structures being formed that were not present in the original image.

None of the diffusion methods produced satisfactory results, so the conclusion was made that more than one image of the same area would be necessary to reduce speckle without compromising the image in other ways.

When more than one image is available, they can be averaged to reduce noise, in particular the effect of speckle. This was investigated in chapter 6. Since minor movement in the head lead to varying shifts in the A-scans, a robust registration of the images is necessary before averaging. The vertical and horizontal registration is done independently. This is possible when horizontal layers, such as the ones in the retina, are present in the image. The vertical registration is based on finding a shortest path by use of regularized dynamic programming through an energyspace based on correlation. A constant horizontal shift is assumed for each image. This one value is found by maximizing correlation.

The averaging method was tested with different settings, but it only had minor effects on the resulting images. This is taken as a sign of the robustness of the method. When combining several images, a significant reduction in speckle is seen, which can make details visible that were not noticeable in the original images. The reduction in noise seemed to wear off when 8-10 images were included.

10.2 Applications

To extract the relevant descriptors in a preoperative image, as was done in chapter 7, it was necessary to find the location of the top of the neuroretina and the top of the hyperreflective layer consisting of the IS/OS and RPE. The last one is found automatically by regularized dynamic programming in the gradient image. The surface of the neuroretina is found semi-automatically with a snake locked to the sides of the image. A few points lying above the surface is provided by the user to ensure the necessary robustness of the method. From the knowledge of the position of these layers, all relevant hole descriptors can be estimated. These are, the center height of the hole, minimum and maximum width of the hole, the hole area and volume and the retinal area.

When looking at a postoperative image, where the macular hole has closed, as was the case in chapter 8, the location of the same two layers are of interest, but also the bottom of the hyperreflective layer. Since the hole has closed, all three can be found automatically with regularized dynamic programming. The interesting descriptors are in this case the thickness of the hyperreflective layer, the thickness of the neuroretina, the width of a possible atrophy, the retinal area and the intensity of a marked section of the outer nuclear layer.

If several OCT scans have been taken at different retinal locations, it is possible to interpolate the neuroretinal thickness found from each scan into an approximate surface. Examples of this is shown in chapter 9. When the data is represented in this way, a swelling or thinning in the neuroretina is visualized in an intuitive way. The shape of a macular hole can not be represented in this way without loosing some of its shape information. But if the neuroretina instead is reconstructed as a 3D surface, this information is maintained, and it is possible to visualize a 3D cup representation of the macular hole.

10.3 Conclusion

OCT is still a technique in development. Alterations such as new lightsources can improve the produced image significantly. These improvements are often the result of new technology being used, at the cost of the final price of the product. If a method based solely on image processing can produce a next generation resulting image from images taken from what is currently a standard apparatus, it would be of great interest.

One way this seems to be achievable is by reducing the effect of speckle in OCT images. In this thesis three different types of diffusion have been applied to OCT images. None of the diffusion methods produced satisfactory results, so an iterative method was developed that averaged several images. Each image is aligned and registered vertically and horizontally to a template, before averaging any images. The method is robust to parametrical changes, and the average image has significantly less noise than the originals.

The extraction of relevant descriptors from pre- and postoperative OCT images of retinas with a macular hole have also been examined. The descriptors can be extracted from the location of transitional layers. They are found automatically or semi-automatically. If these layers are known for several slices located at different retinal positions, the neuroretinal thickness can be represented as a surface map or 3D surface, in this way visualizing the entire retina instead of slices of it. The software developed throughout the course of this project is to be used in a case study at Herlev Hospital, where different surgical techniques to treat macular hole are evaluated. The study will hopefully lead to new insight about the optimal treatment and pathogenesis of a macular hole.

Part V

Appendix

Appendix A

Matlab Code Overview

Relevant code from chapter 5 to 9 has been supplied on the enclosed CD. The scripts are briefly described. For more information see the appropriate file. If a script uses any functions it is mentioned.

Chapter 4

- **snake=DynProgReg(X, lambda, p, RescaleType, q)**
  Regularized dynamic programming on energyspace X

Chapter 5

- **imout=AlignImByCorr(IM)**
  Aligns image such that the correlation between columns are maximized
  uses: shift
- **imout=LinearDiffusion(IM, sigma)**
  Linear diffusion by gaussian filtering
- **imout=ComplexDiffusion(IM, ang, k, mu, N)**
  Ramp preserving complex diffusion

- **imout=CoherenceDiffusion(IM, sigma1, sigma2, fractileK, N)**
  Coherence enhancing diffusion
  uses: makeColorMap
- **out=shift(in, nbr)**
  shifts column according to nbr
- **makeColorMap**
  Makes colormap for orientation image

Chapter 6

- **AverageRunner**
  Aligns and registers images and produces average image
  uses: DynProgReg, AlignImByCorr and shift

Chapter 7

- **FindSurfacesMH**
  Locates center and top of hyperreflective layer automatically and finds top of retina with active contour initialized from points provided by user
  uses: DynProgReg, shift, ParamSnake, FreeHandSnake and DrawSnake
- **DetermineHoleParameters**
  Determines descriptors for macular hole
  uses: DrawSnake and shift
- **v=ParamSnake(IMres, v, iterations, alpha, beta, delta, gamma)**
  Active contour, where endpoints are kept on sides of image
  uses: balloon, parameterize and DrawSnake
- **DrawSnake(v, color)**
  Draws snake on figure
- **v=FreeHandSnake(go)**
  Get points from user by clicking on figure
- **v=parameterize(x, y, N)**
  Returns N equally spaced points on line through points
- **Fb=balloon(v)**
  Returns inwards normal to every points
- **[v2 sizev2]=interpolatePoints(v, nr)**
  Interpolates v such that all points returned are connected

Chapter 8

- **FindRPEDynProg**
  Locates top and bottom of hyperreflective layer
  uses: DynProgReg
- **FindRetinaTop**
  Locates top of retina and calculate descriptors
  uses: DynProgReg

Chapter 9

- **MapSurface**
  Maps surface from six radial scans
  uses: convertVto3D
- **prepRBF**
  Generates 3D points for RBFtest
  uses: DrawSnake, balloon, parameterize
- **RBFtest**
  Generates 3D surface by radial basis functions
  uses: FastRBF Toolbox
- **[x y z]=convertVto3D(v)**
  Returns 3D coordinates from six radial scans

Appendix B

Test of Averaging Method

The method to reduce noise by averaging several images presented in chapter 6 has been tested on 7 sets of OCT images. From each set between 10 and 13 images have been used to produce the final image. One set is of a healthy individual, 3 sets have a preoperative macular hole and 3 sets are postoperative images of a macular hole. One initial image, an aligned version of the initial image and the average image for each set are shown in this chapter.

For a general discussion of the images see section 6.3.

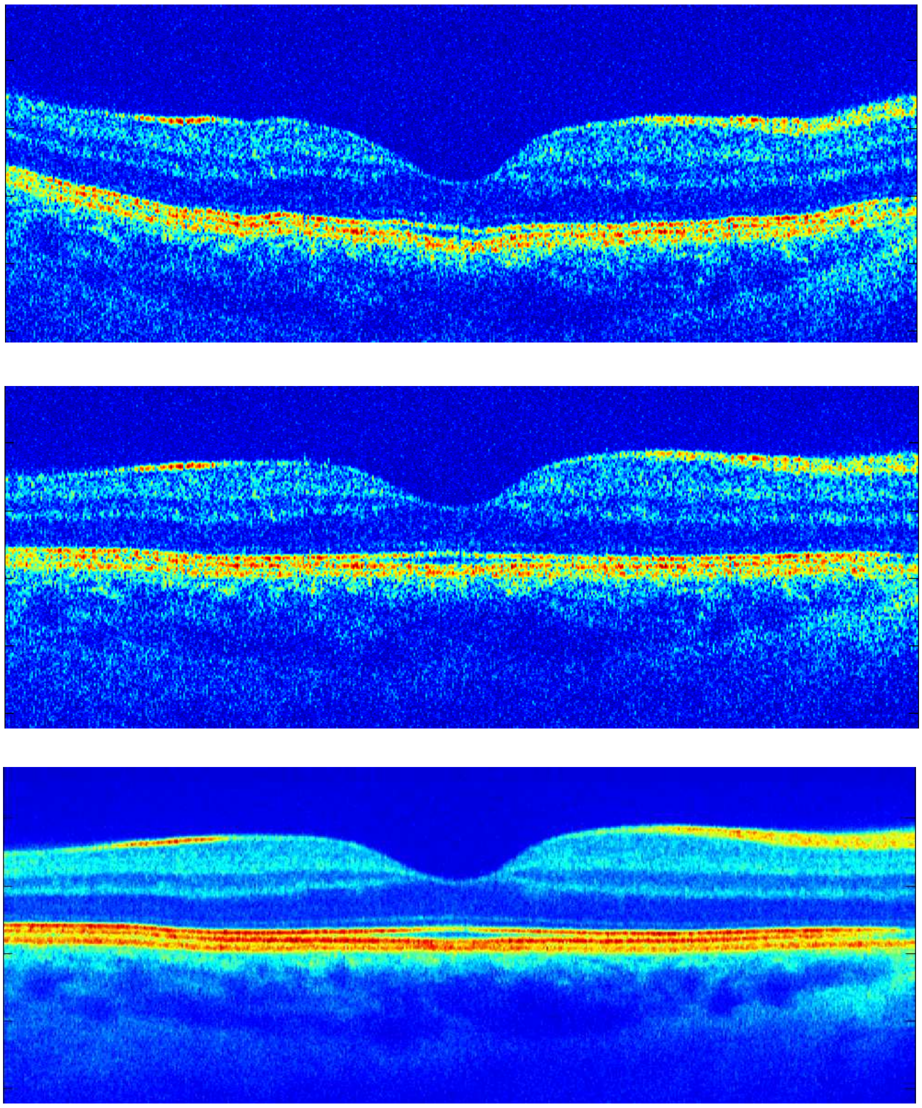

Figure B.1: Healthy individual.

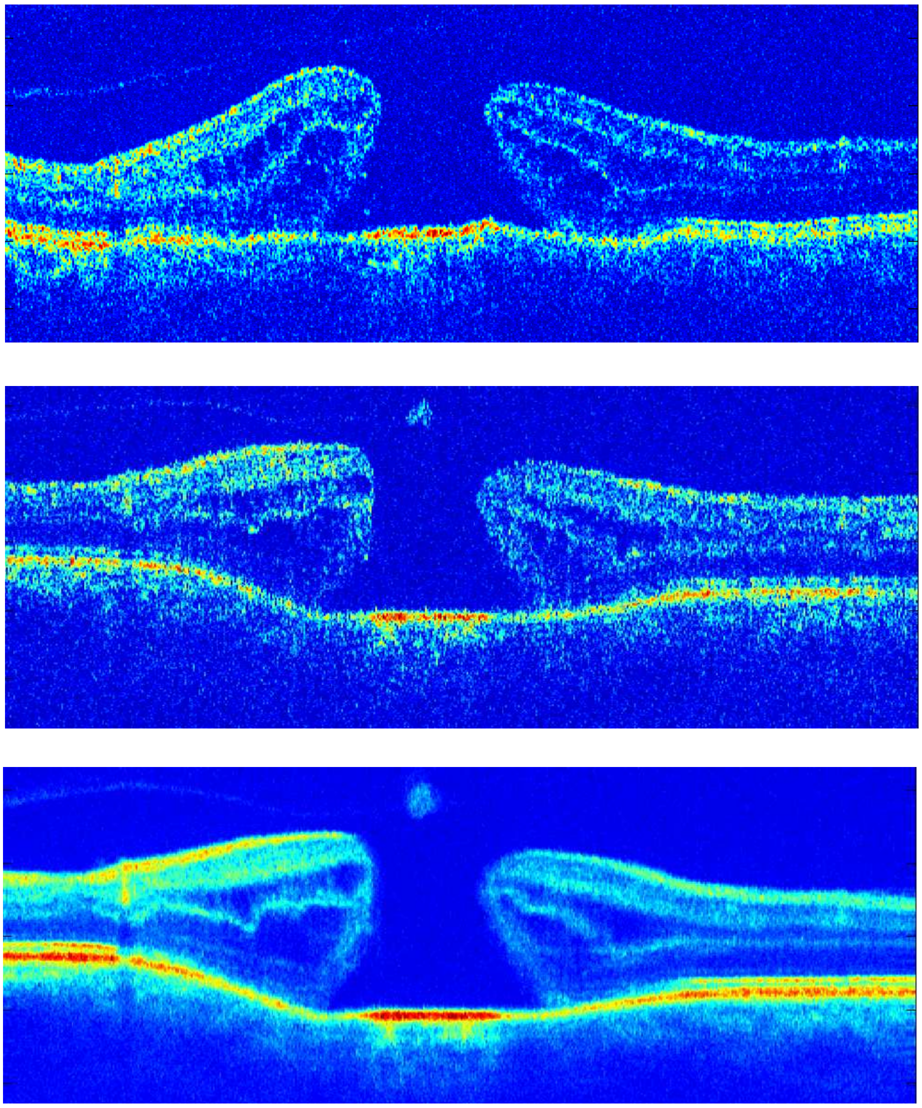

Figure B.2: Macular hole set 1.

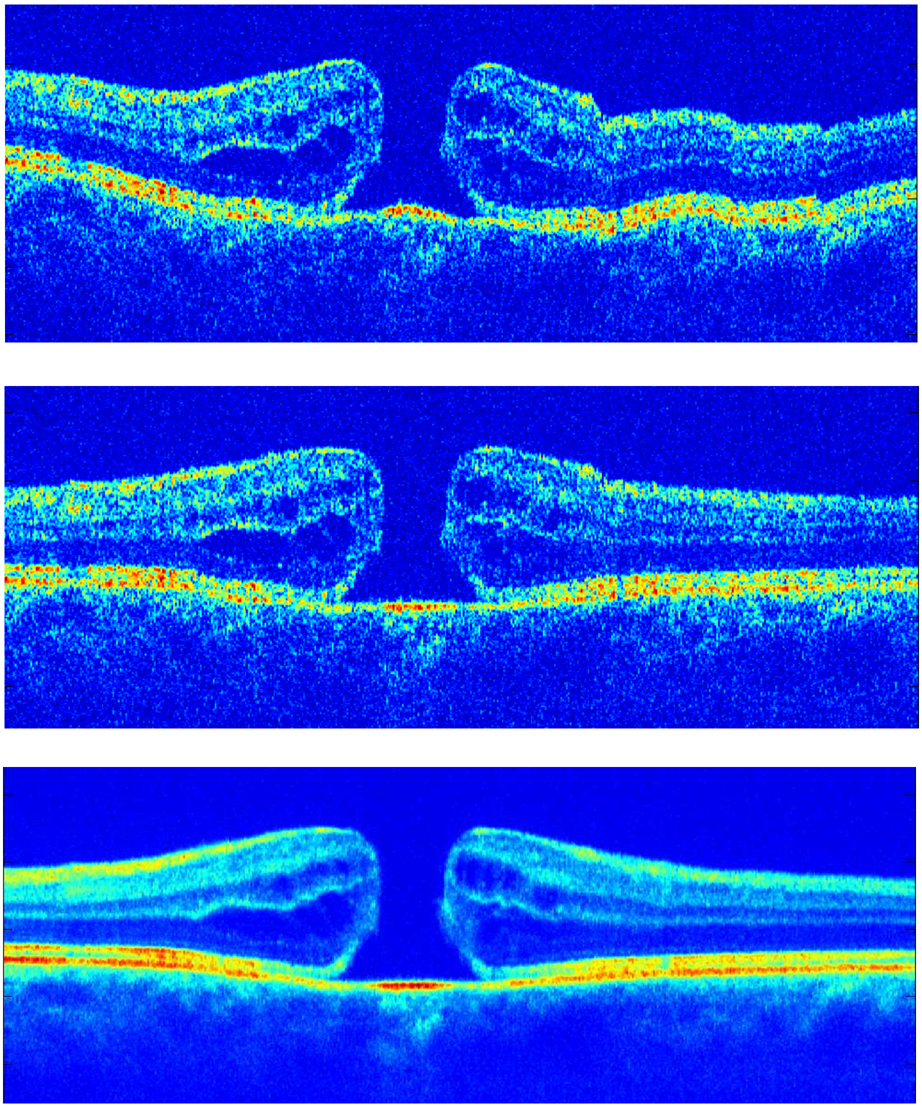

Figure B.3: Macular hole set 2.

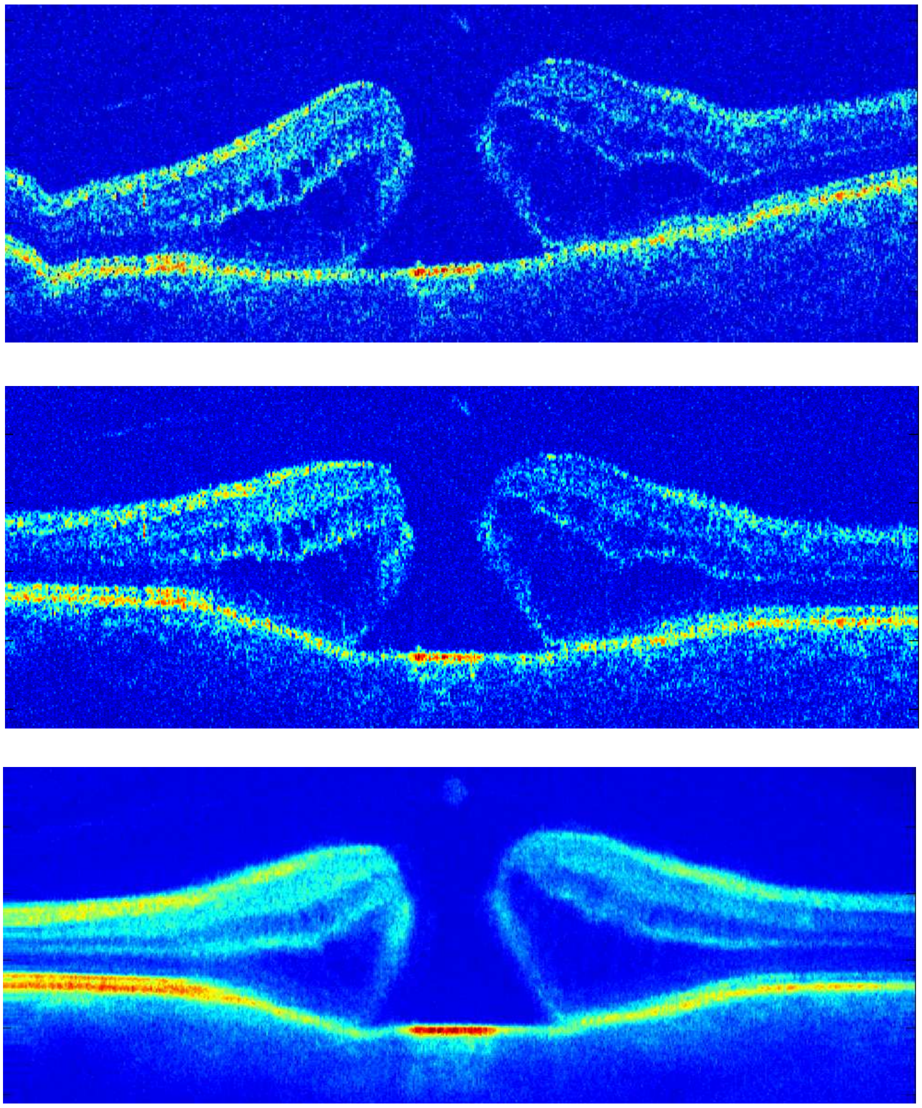

Figure B.4: Macular hole set 3.

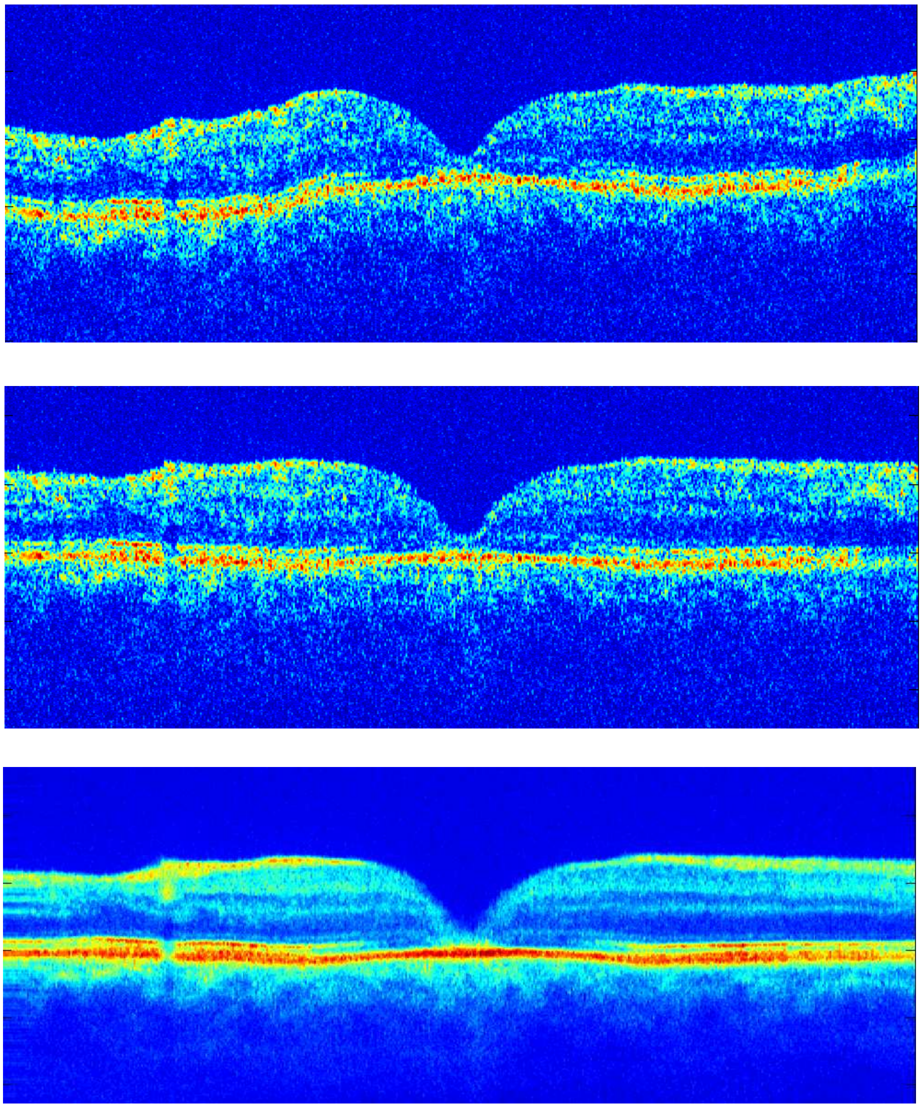

Figure B.5: Postoperative macular hole set 1.

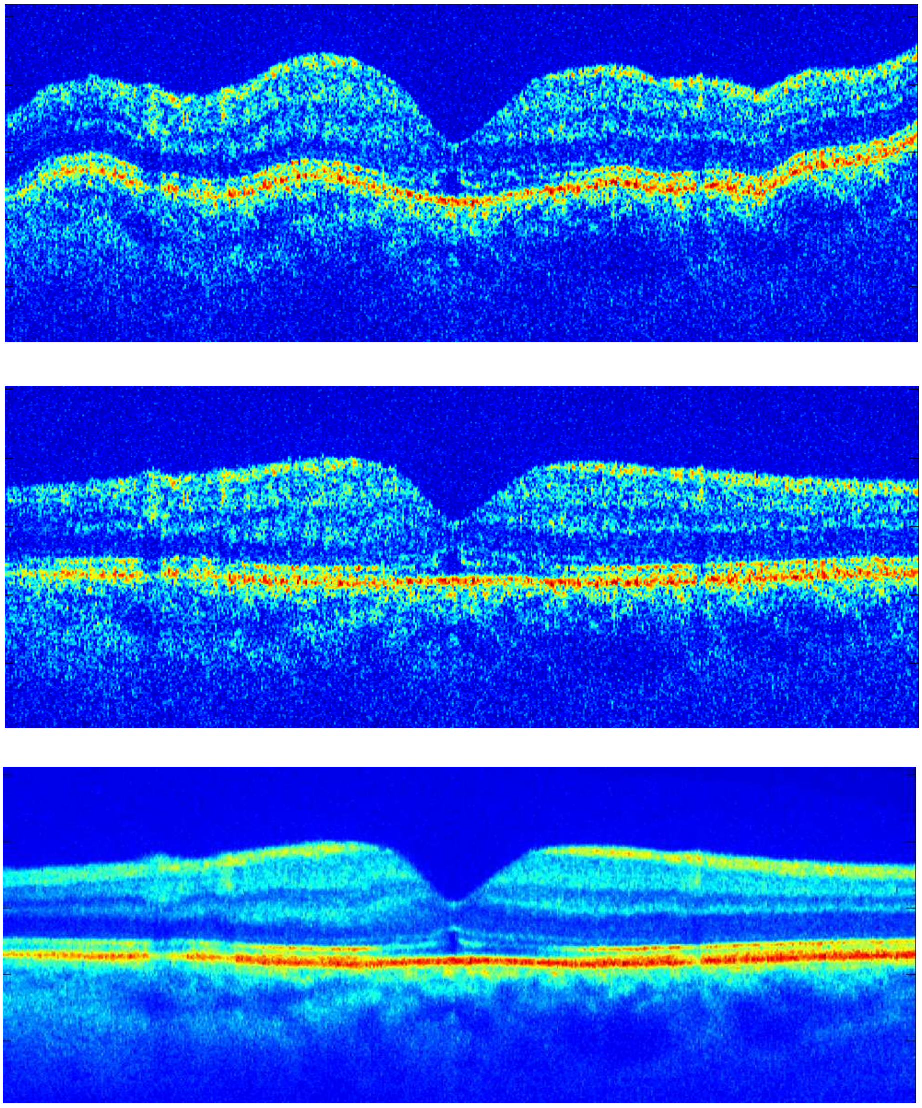

Figure B.6: Postoperative macular hole set 2.

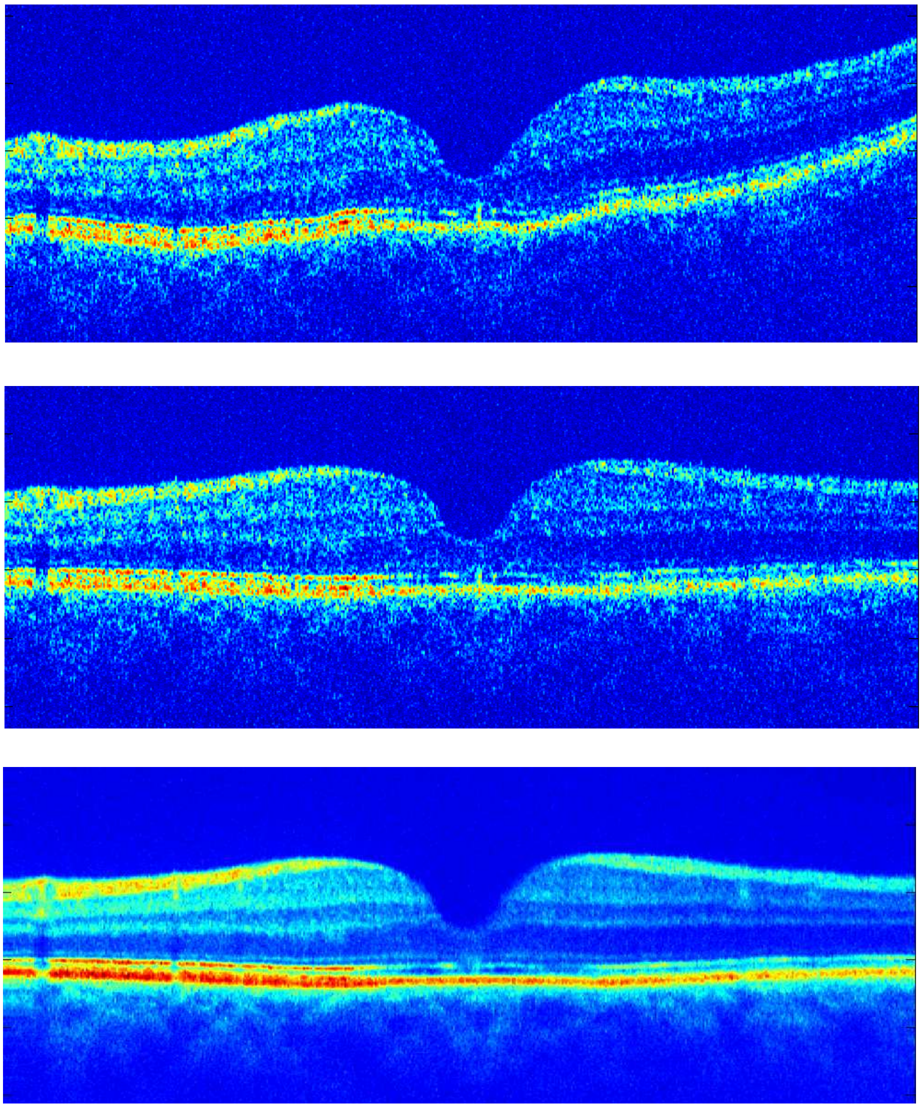

Figure B.7: Postoperative macular hole set 3.

Appendix C

Registration Failure

In the following an imageset consisting of 11 images taken of the same retinal location is used to show a registration failure with a less robust method. One of the original images are shown in figure C.1. On the left side of the macula the RNFL and the RPE are very distinct but on the right side there is a pathology that affects the retina.

The left side is handled well by a method that maximizes the correlation for every A-scan, but it can not handle the right side, where wrong vertical shifts are estimated as seen in C.2. Horizontal discontinuities occur in the structure of the image. With a regularized dynamic programming method in energyspace the right side of the image is registered much better, shown in figure C.3. The method gives no discontinuities in the horizontal structure due to the continuity constraint in energyspace.

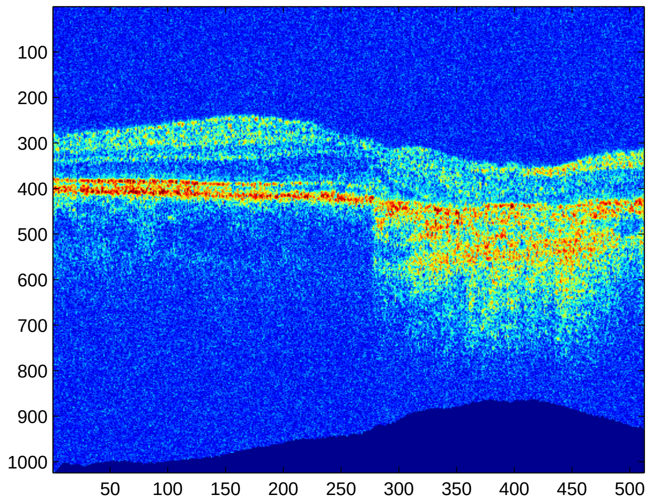

Figure C.1: One image from an imageset consisting of 11 images taken of the same retinal area. On the left side of the macula the RNFL and the RPE are very distinct but on the right side there is a pathology that affects the retina.

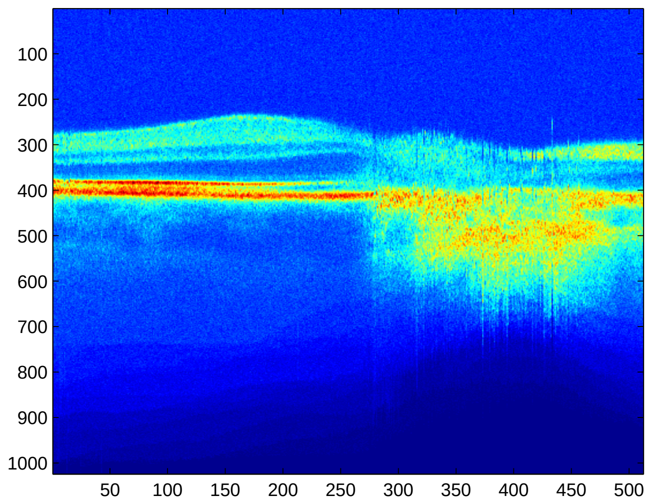

Figure C.2: A method that maximizes the correlation for every A-scan has been used to register an imageset. The left side is handled well, but it can not handle the right side, where wrong vertical shifts are estimated.

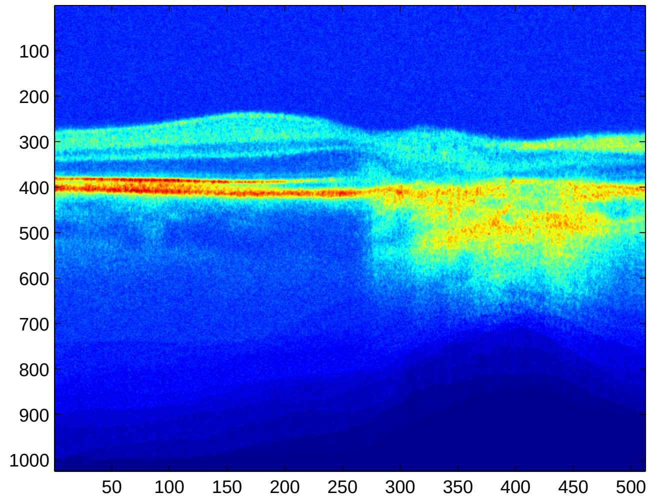

Figure C.3: Regularized dynamic programming in energyspace has been used to register the imageset. the right side of the image is registered better than the simple implementation, shown in figure C.2. The method gives no discontinuities in the horizontal structure due to the continuity constraint in energyspace.

Appendix D

Test of Pre-OMaH Tool

The Pre-OMaH Tool has been tested on six very different macular hole cases. The images have not been chosen for their high quality, they should represent typical images that have not been averaged. The robustness of finding the top of the hyperreflective layer and in particular the top of the retina is to be tested. For all the images, between 12 and 17 points have been given by the user just outside the retina to initialize the active contour.

For all 6 test images, the top of the hyperreflective layer only needed assistance in case 3. Locating the top of the retina was harder as expected. The top of the retina on test images 1-3 have been found without significant adjustment on the force parameters, but test image 4 and 5 had so narrow macular holes that the bending force parameter needed to be reduced significantly to locate the borders of the hole. The same was necessary for test image 6, where the lid of the hole is still attached on the right side of the image. The contrast is so poor in test images 5 and 6 that the balloon force also needed to be reduced significantly, or else the snake would not be stopped by the borders of the retina.

For some of the images a significant parameter tuning was required to locate the top of the retina. This was because of narrow macular holes or a low contrast. Never the less, satisfactory solutions have be found for all images. The contrast of the images can be increased when using an average image instead of a single image, thus reducing the time spend on adjusting parameters, and hopefully also finding an even better solution.

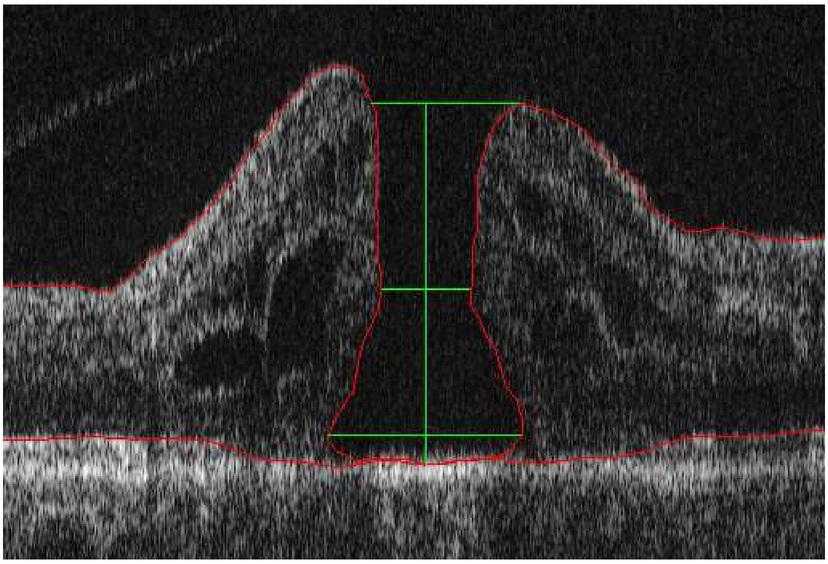

(a) Preoperative test image 1

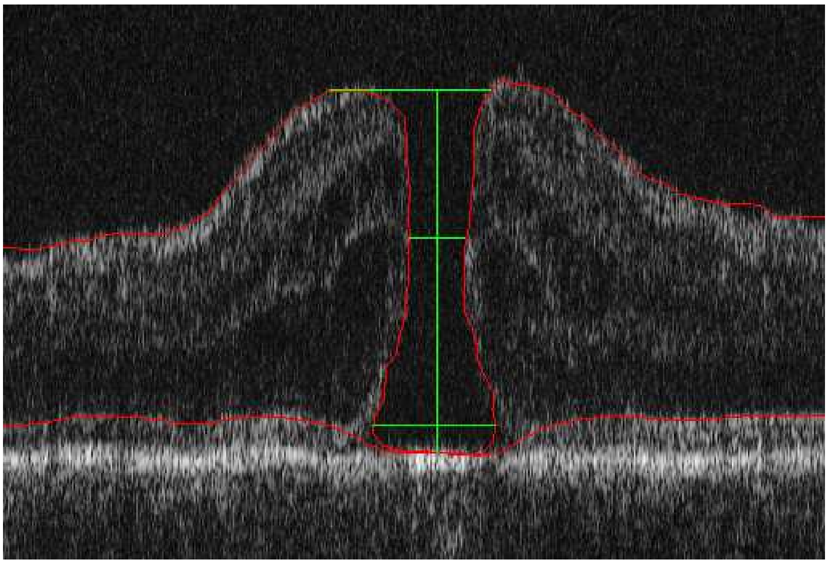

(b) Preoperative test image 2

Figure D.1: Preoperative test image 1 and 2.

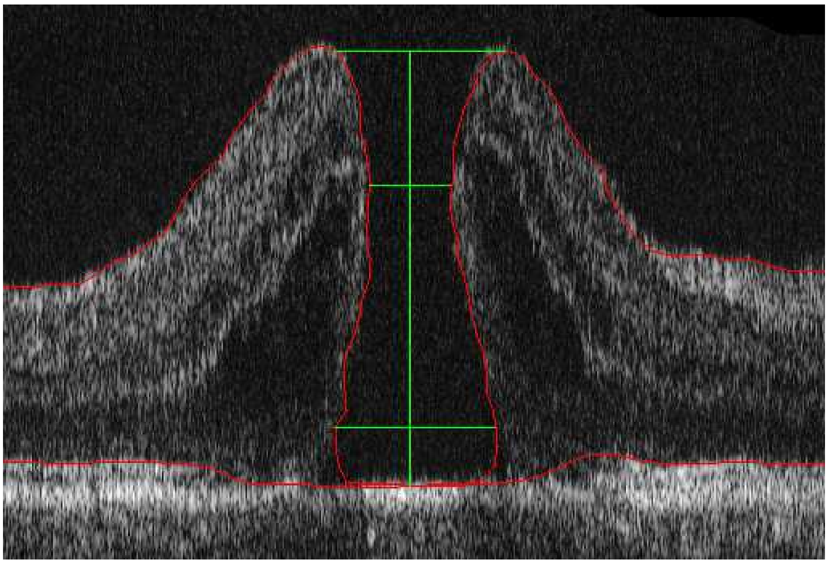

(a) Preoperative test image 3

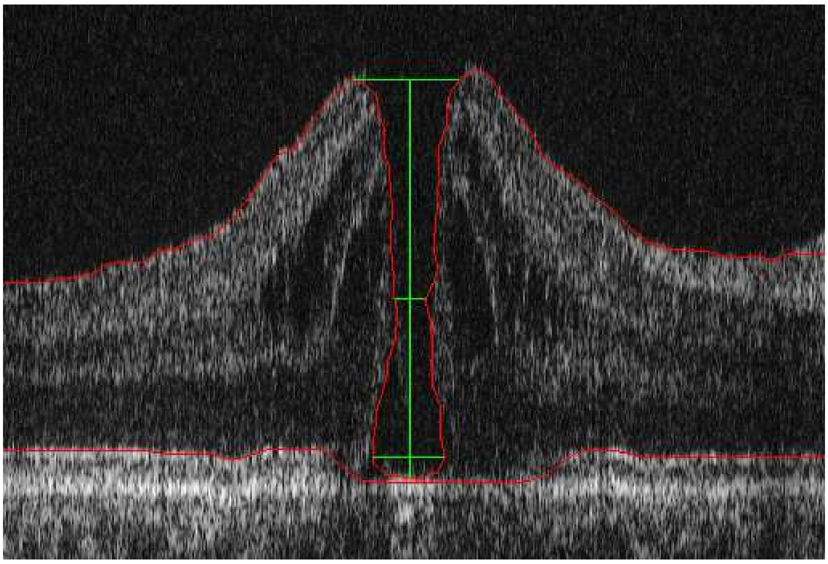

(b) Preoperative test image 4

Figure D.2: Preoperative test image 3 and 4.

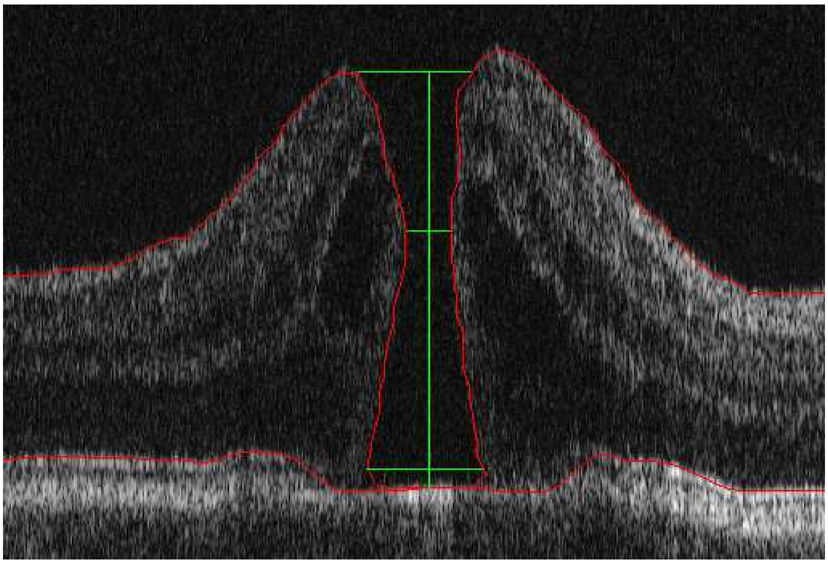

(a) Preoperative test image 5

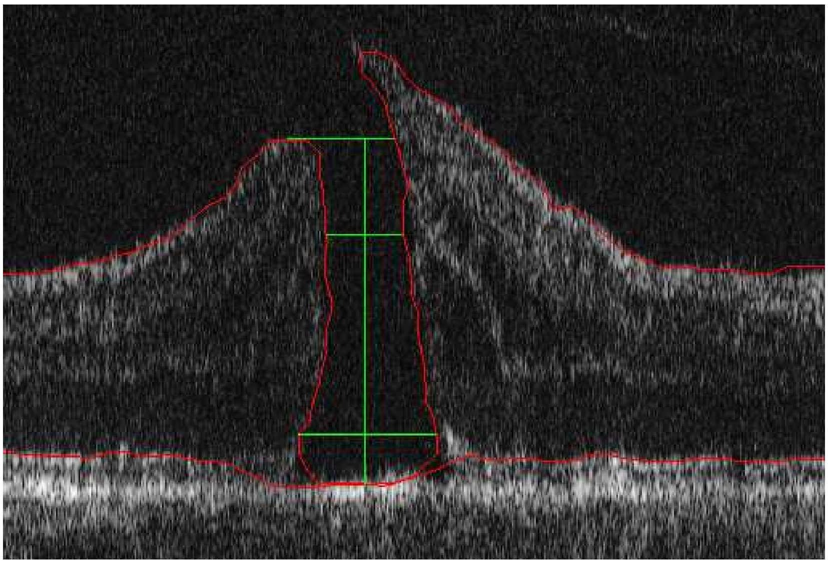

(b) Preoperative test image 6

Figure D.3: Preoperative test image 5 and 6.

Appendix E

Test of Post-OMaH Tool

To test Post-OMaH Tool, six different postoperative cases have been used. No averaging of images have been done prior to analysis. The aim of this test is not to try to evaluate whether or not an atrophy is present that is up to the experts to decide. Instead it is to find out how much assistance is needed to locate the three layers.

The layers seen on the test images 1-3 have been located automatically without any assistance. In test image 4 the bottom of the RPE was not correctly found, it needed a few points to pull the path closer to the actual location. In test image 5 the top of the hyperreflective layer, was on the left side located at the top of the RPE. This was corrected by a few of points. Test image 6 was chosen for its lack of contrast on the right side, and the top of the hyperreflective layer fell to the top of the RPE as expected. The top of the retina also needed assistance in the far right side.

In general the algorithms succeeds when the image has a high contrast at the transitions that should be located, in particular between the RPE and choroid. But even with a severely degraded image as test image 6, a qualified initial guess was given that could be pulled toward the preferred location with a limited set of points provided by the user.

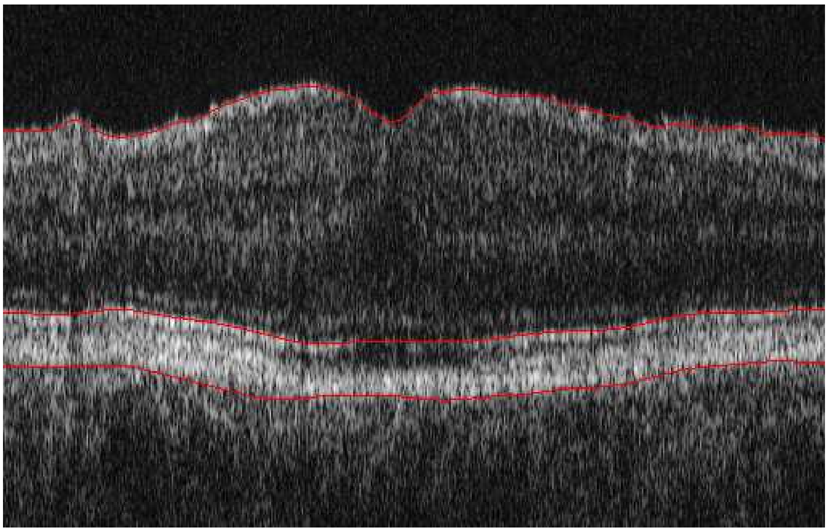

(a) Postoperative test image 1

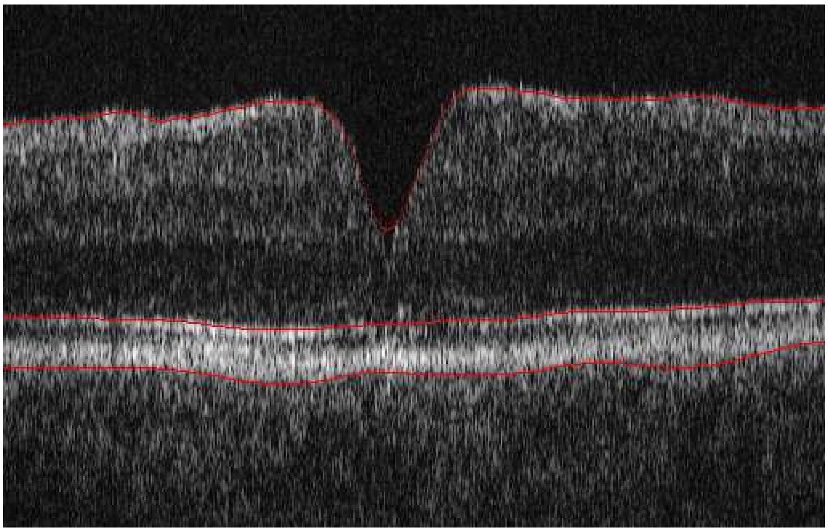

(b) Postoperative test image 2

Figure E.1: Postoperative test image 1 and 2.

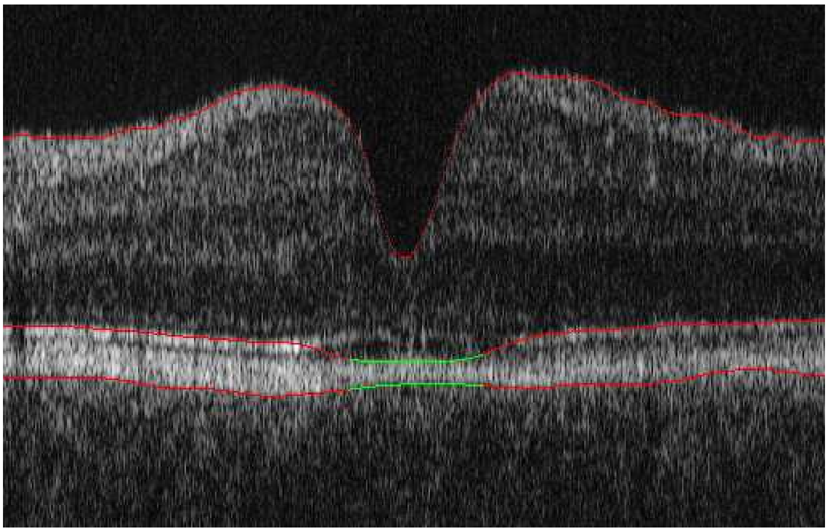

(a) Postoperative test image 3

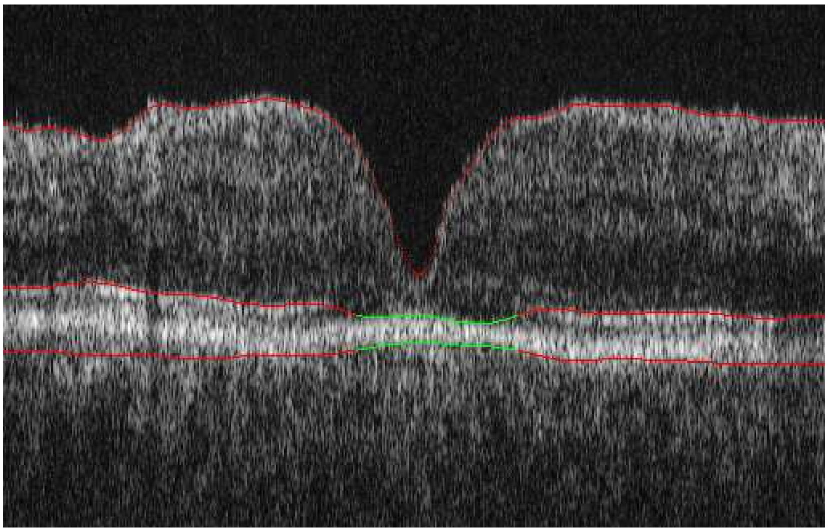

(b) Postoperative test image 4

Figure E.2: Postoperative test image 3 and 4.

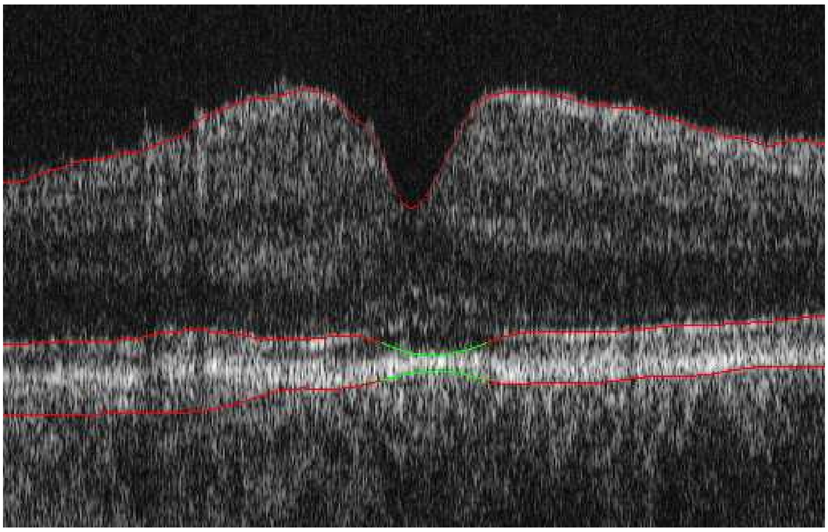

(a) Postoperative test image 5

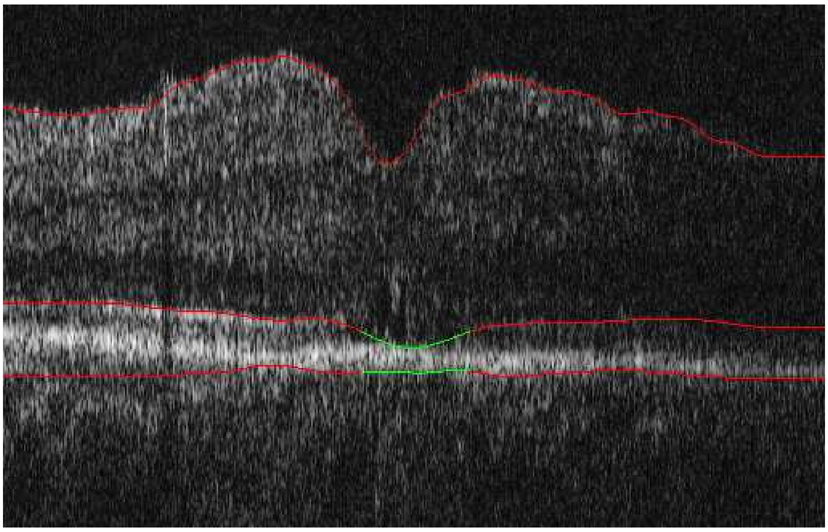

(b) Postoperative test image 6

Figure E.3: Postoperative test image 5 and 6.

Appendix F

User manual for Pre-OMaH Tool

In order to be able to determine the relevant preoperative descriptors, the top of the hyperreflective layer and the the retina must first be located. The process implemented in Pre-OMaH Tool will be described in the following. For a screenshot of Pre-OMaH Tool see figure 7.7.

- **Loading an image**

  Before loading an image, choose whether or not the loaded image should be aligned, and whether this should be done to the RPE, or simply by correlating the A-scans, by un/checking the boxes. Two types of image can be loaded, either a .bmp grey-scale image or a .raw exported from StratusOCT.

- **Hyperreflective layer**

  Once the image is shown, choose how much postoperative smoothing should be done on the path found as the the top of the hyperreflective layer. Press "Calculate", and the top of the hyperreflective layer is located. You will be prompted whether or not you are satisfied with the result. If you are not, provide points for problematic areas, and press "Calculate". This process can be repeated, until the result is satisfactory.

- **Top of retina**

  Provide an initial path lying just outside the actual surface, by clicking on the image, starting from the left. The first and last point will be connected with a horizontal line to the sides. Press "Calculate", and evaluate whether or not the result is satisfactory. If not adjust the three parameters appropriately, and provide a new initial path. Once the top is found correctly, all descriptors are calculated and marked on the image. They are explained below. The image with lines drawn on can be saved, and the descriptors can be saved to a text file.

Height - Height of the macular hole in the center in μm.

Min. Width - The minimum width of the macular hole in μm.

Max. Width - The maximum width of the macular hole in μm. The maximum widht must lie below the minimum width.

Area of retina - Area lying between the top of the retina and the top of the hyperreflective layer located in mm^2.

Area of hole - Area of the hole in mm^2, where the top is defined as where liquid poured into the hole would overflow.

Volume of hole - The volume in mm^3 of the object produced by rotating the macular hole around the center line.

Appendix G

User manual for Post-OMaH Tool

To determine the relevant postoperative descriptors, the different layers must first be located, along with where the ONL should be sampled. The process implemented in Post-OMaH Tool will be described in the following. For a screenshot of Post-OMaH Tool see figure 8.8.

- **Loading an image**

  Before loading an image, choose whether or not the loaded image should be aligned, by un/checking the box. Two types of image can be loaded, either a .bmp grey-scale image or a .raw exported from StratusOCT.

- **Hyperreflective layers**

  Once the image is shown, choose how much postoperative smoothing should be done on the three paths found. Press "Calculate", and the top and bottom of the hyperreflective layer is located. You will be prompted whether or not you are satisfied with the result. If you are not, provide points for problematic areas, first top, then center and bottom of hyperreflective layer, pressing "Calculate" between each. This process can be repeated, until the result is satisfactory.

- **Top of retina**

  The top of the retina is then found, and points can be provided in the same way as for the hyperreflective layer, if the result is not satisfactory.

- **Sampling the ONL**

  Once the top is found correctly, the areas to calculate the center and side intensity of the ONL should be chosen by selecting two corners in the rectangle in the center and two corners in the rectangle on the side. All descriptors are then calculated. They are explained below. The image with lines drawn on can be saved, and the descriptors can be saved to a text file.

Retinal Thickness - Neuroretinal thickness in μm, sampled at predefined locations ranging from 2mm to one side to 2mm on the other side.

Hyperref. layer - Thickness of the hyperreflective layer in μm, sampled at predefined locations ranging from 2mm to one side to 2mm on the other side.

Intensities - Relative intensity of the two selected rectangles in the fovea and on the side, and the ratio of the two.

Std - Standard deviation of the pixels selected to be sampled for the ONL.

Atrophy width - Width of the estimated atrophy in μm, marked in green on the image.

Retinal Area - Area lying between the top of the retina and the top of the hyperreflective layer located in mm^2.

Bibliography

[1] Jakob Friis Morten la Cour. Macular holes: classification, epidemiology, natural history and treatment. *Acta Ophthalmologica Scandinavica*, 80(6):579–587, 2002.

[2] SA. Madreperla and BW. II McCuen. *Macular Hole. Pathogenersis, Diagnosis and Treatment*, chapter Internal limiting membrane removal in surgery for full-thciness macular holes, pages 125–146. Boston: Butterworth Heinemann, 1999.

[3] H. Logan Brooks. Macular hole surgery with and without internal limiting membrane peeling. *Ophthalmology*, 107(10):1939–1948, 2000.

[4] David Huang, Eric A. Swanson, Charles P. Lin, Joel S. Schuman, William G. Stinson, Warren Chang, Michael R. Hee, Thomas Flotte, Kenton Gregory, Carmen A. Puliafto, and James G. Fujimoto. Optical coherence tomography. *Science*, 254(5035):1178–1181, 1991.

[5] Thomas Martini Jorgensen, Bjarne Ersboll, Birgit Sander, and Michael Larsen. Reducing speckle noise in retinal oct images by aligning multiple b-scans. *Progress in Biomedical Optics and Imaging - Coherence Domain Optical Methods and Optical Coherence Tomography in Biomedicine VIII and Proceedings of SPIE - The International Society for Optical Engineering*, 5316:205–213, 2004.

[6] J.M. Schmitt. Optical coherence tomography (oct): a review. *Selected Topics in Quantum Electronics, IEEE Journal on*, 5(4):1205 –1215, 1999.

[7] M. R. Hee, J. A. Izatt, E. A. Swanson, D. Huang, J. S. Schuman, C. P. Lin, C. A. Puliafito, and J. G. Fujimoto. Optical coherence tomography of the human retina. *Arch Ophthalmol*, 113(3):325–332, 1995.

[8] Egbert Lenderink. Optical coherence tomography for three-dimensional imaging of skin features. *Proceedings of SPIE - The International Society for Optical Engineering*, 3567:70–77, 1999.

[9] C.K. Hitzenberger, P. Trost, Pak-Wai Lo, and Qienyuan Zhou. Three-dimensional optical coherence tomography of the human retina in vivo by high-speed transversal scanning. *Proceedings of the SPIE - The International Society for Optical Engineering*, 5316(1):7–11, 2004.

[10] Oregon medical laser center, http://omlc.ogi.edu/news/jan98/gallery_jan98.html.

[11] J.M. Schmitt, S.H. Xiang, and K.M. Yung. Speckle in optical coherence tomography: an overview. *Proceedings of SPIE - The International Society for Optical Engineering*, 3726:450–461, 1999.

[12] Yuan An, JianQuan Yao, and RuiKang Wang. Analysis of speckle in optical coherence tomography. *Proceedings of SPIE - The International Society for Optical Engineering*, 4916:245–250, 2002.

[13] Stephane Paes, Seon Young Ryu, Jihoon Na, EunSeo Choi, Changsu Lee, and Byeong Ha Lee. Combined applications of iterative deconvolution methods and adaptive speckle filters for optical coherence tomography. *Biophotonics, 2004. APBP 2004. The Second Asian and Pacific Rim Symposium on*, pages 137–138, 2004.

[14] H. Choi, T.E. Milner, and A.C. Bovik. Speckle noise reduction and segmentation on polarization sensitive optical coherence tomography images. *Engineering in Medicine and Biology Society, 2003. Proceedings of the 25th Annual International Conference of the IEEE*, pages 1062–1065 Vol.2, 2003.

[15] Desmond C. Adler, Tony H. Ko, and James G. Fujimoto. Speckle reduction in optical coherence tomography images by use of a spatially adaptive wavelet filter. *Optics Letters*, 29(24):2878–2880, 2004.

[16] J. Rogowska and M.E. Brezinski. Evaluation of the adaptive speckle suppression filter for coronary optical coherence tomography imaging. *Medical Imaging, IEEE Transactions on*, 19(12):1261 –1266, 2000.

[17] D.C. Fernandez. Delineating fluid-filled region boundaries in optical coherence tomography images of the retina. *Medical Imaging, IEEE Transactions on*, 24(8):929–945, 2005.

[18] Dr. john grohol's psych central, http://psychcentral.com/psypsych/eye.

[19] B Sander, M Larsen, L Thrane, J L Hougaard, and T M Jorgensen. Enhanced optical coherence tomography imaging by multiple scan averaging. *British Journal of Ophthalmology*, 89(2):207–212, 2005.

[20] JD Gass. Idiopathic senile macular hole: its early stages and pathogenesis. *Arch Ophthalmol*, 106:629–639, 1988.

[21] W.E. Smiddy and H.W. Flynn. Pathogenesis of macular holes and therapeutic implications. *American Journal of Ophthalmology*, 137(3):525–537, 2004.

[22] J. Wender, T. Iida, and L.V. Del Priore. Morphologic analysis of stage 3 and stage 4 macular holes: Implications for treatment. *American Journal of Ophthalmology*, 139(1):1–10, 2005.

[23] R E Bellman. *Dynamic Programming.* Princeton University Press, 1957.

[24] A.A. Amini, T.E. Weymouth, and R.C. Jain. Using dynamic programming for solving variational problems in vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(9):855–867, 1990.

[25] Michael Buckley and Jean Yang. Regularised shortest-path extraction. *Pattern Recognition Letters*, 18(7):621–629, 1997.

[26] Thomas Martini Jorgensen, Jakob Thomadsen, Lars Thrane, and Peter E. Andersen. Dynamic contour model for aligning and segmenting oct b-scans. *Coherence Domain Optical Methods and Optical Coherence Tomography in Biomedicine IX and Progress in Biomedical Optics and Imaging - Proceedings of SPIE*, 5690:480–485, 2005.

[27] K.J. Coakley and P. Hale. Alignment of noisy signals. *Instrumentation and Measurement, IEEE Transactions on*, 50(1):141 –149, 2001.

[28] R. Jane, H. Rix, P. Caminal, and P. Laguna. Alignment methods for averaging of high-resolution cardiac signals: a comparative study of performance. *Biomedical Engineering, IEEE Transactions on*, 38(6):571–579, 1991.

[29] G. Gilboa, N. Sochen, and Y.Y. Zeevi. Image enhancement and denoising by complex diffusion processes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(8):1020–1036, 2004.

[30] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(7):629 –639, 1990.

[31] Xiaohui Hao, Shangkai Gao, and Xiaorong Gao. A novel multiscale nonlinear thresholding method for ultrasonic speckle suppressing. *Medical Imaging, IEEE Transactions on*, 18(9):787 –794, 1999.

[32] A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy. A versatile wavelet domain noise filtration technique for medical imaging. *Medical Imaging, IEEE Transactions on*, 22(3):323–331, 2003.

[33] Joachim Weickert. Coherence-enhancing diffusion of colour images. *Image and Vision Computing*, 17(3-4):201–212, 1999.

[34] F. Sattar, L. Floreby, G. Salomonsson, and B. Lovstrom. Image enhancement based on a nonlinear multiscale method. *Image Processing, IEEE Transactions on*, 6(6):888–895, 1997.

[35] Xiaomei Song, Brian W. Pogue, Shudong Jiang, Marvin M. Doyley, Hamid Dehghani, Tor D. Tosteson, and Keith D. Paulsen. Automated region detection based on the contrast-to-noise ratio in near-infrared tomography. *Applied Optics*, 43(5):1053–1062, 2004.

[36] S. Kusuhara, M.F. Teraoka Escano, S. Fujii, Y. Nakanishi, Y. Tamura, A. Nagai, H. Yamamoto, Y. Tsukahara, and A. Negi. Prediction of post-operative visual outcome based on hole configuration by optical coherence tomography in eyes with idiopathic macular holes. *American Journal of Ophthalmology*, 138(5):709–716, 2004.

[37] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: active contour models. *International Journal of Computer Vision*, 1(4):321–31, 1987.

[38] N. Villate, J.E. Lee, A. Venkatraman, and W.E. Smiddy. Photoreceptor layer features in eyes with closed macular holes: Optical coherence tomography findings and correlation with visual outcomes. *American Journal of Ophthalmology*, 139(2):280–289, 2005.

[39] FarField Technology. *FastRBF Toolbox*, 2004.