
Theses and Dissertations

Fall 2013

A combined machine-learning and graph-based framework for the 3-D automated segmentation of retinal structures in SD-OCT images

Bhavna Josephine Antony
University of Iowa

Follow this and additional works at: <https://ir.uiowa.edu/etd>



Part of the [Electrical and Computer Engineering Commons](#)

Copyright 2013 Bhavna Josephine Antony

This dissertation is available at Iowa Research Online: <https://ir.uiowa.edu/etd/4944>

Recommended Citation

Antony, Bhavna Josephine. "A combined machine-learning and graph-based framework for the 3-D automated segmentation of retinal structures in SD-OCT images." PhD (Doctor of Philosophy) thesis, University of Iowa, 2013.

<https://doi.org/10.17077/etd.ny1b4j7x>

Follow this and additional works at: <https://ir.uiowa.edu/etd>



Part of the [Electrical and Computer Engineering Commons](#)

A COMBINED MACHINE-LEARNING AND GRAPH-BASED FRAMEWORK
FOR THE 3-D AUTOMATED SEGMENTATION OF RETINAL STRUCTURES
IN SD-OCT IMAGES

by

Bhavna Josephine Antony

A thesis submitted in partial fulfillment of the
requirements for the Doctor of Philosophy degree
in Electrical and Computer Engineering
in the Graduate College of
The University of Iowa

December 2013

Thesis Supervisor: Assistant Professor Mona K. Garvin

Copyright by
BHAVNA JOSEPHINE ANTONY
2013
All Rights Reserved

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Bhavna Josephine Antony

has been approved by the Examining Committee for the thesis requirement for the Doctor of Philosophy degree in Electrical and Computer Engineering at the December 2013 graduation.

Thesis Committee: _____

Mona K. Garvin, Thesis Supervisor

Michael D. Abramoff

Milan Sonka

Joseph M. Reinhardt

Edward Ratner

ACKNOWLEDGEMENTS

I would like to begin by thanking Dr. Mona Garvin, my advisor, for all of her guidance and support throughout this project. I would also like to thank Dr. Milan Sonka and Dr. Michael Abramoff for all their advice, encouragement and invaluable inputs during my time here at the University of Iowa. I also received considerable assistance from numerous ophthalmologists during my work here. In particular, I would like to thank Drs. Elliott Sohn, Woojin Jeong and Matthew Harper for all the assistance provided with respect to the animal data. A large portion of my thesis work would have been impossible without their collaboration.

I would also like to make particular mention of Dr. Woojin Jeong, Dr. Priya Gupta, Dr. Pavlina Sonkova, Saleh Miri, Dr. Mona Garvin and Dr. Young Kwon for their effort and time spent on tracing structures in SD-OCT images. The creation of reference sets is an extremely time consuming process and I am very grateful to all of them for taking on this extremely important, yet mind-numbing task.

A number of my fellow graduate students and colleagues also deserve special mention. In particular, I would like to thank Kit Lee, Raykai Wang, Li Tang, Zhihong Hu, Qi Song and Junjie Bai for all their help over the years.

I also thank my parents and elder sister for being supportive of this seemingly never-ending educational endeavor. Cathy Kern also deserves special mention here, for willingly stepping into the role of ‘Iowa mom’ and being an endless source of encouragement.

And last, but certainly not least, I’d like to thank Mark Christopher for his encouragement, the intense discussions and willingness to read endless revisions of my thesis. But I’m mostly grateful for his patience (especially when the panic attacks were barely 30 minutes apart!) and all the times he calmly said “You’re going to be fine, Bhavs.”

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	vii
CHAPTER	
1 INTRODUCTION	1
1.1 Thesis Overview	5
2 CLINICAL BACKGROUND	7
2.1 Clinical Motivation	7
2.2 Retinal Structures	8
3 TECHNICAL BACKGROUND	12
3.1 Graph-Theoretic Method	12
3.1.1 Feasibility Constraints	12
3.1.2 Cost Functions	15
3.1.3 Graph Construction	17
3.2 Review of Prior Approaches Proposed for the Segmentation of Retinal Surfaces in SD-OCT Images	17
4 A 3-D AUTOMATED METHOD FOR THE CORRECTION OF AXIAL ARTIFACTS IN SD-OCT VOLUMES	24
4.1 Data	26
4.2 Methods	28
4.3 Experimental Methods	31
4.3.1 Validation using paired OCT scans with orthogonal fast-scanning axes.	32
4.3.2 Validation using 3-D reconstructions of the optic nerve head.	33
4.4 Modeling the Shape of the Eye	35
4.5 Results	36
4.6 Discussion & Conclusion	41
5 INCORPORATION OF TEXTURE-BASED FEATURES IN AN OPTIMAL GRAPH-THEORETIC APPROACH WITH APPLICATION TO THE 3-D SEGMENTATION OF INTRARETINAL SURFACES IN SD-OCT VOLUMES	43
5.1 Methods	45
5.2 Experimental Methods	48

5.3	Results	49
5.4	Discussion	54
6	AUTOMATED 3-D SEGMENTATION OF RETINAL SURFACES IN SD-OCT VOLUMES NORMAL AND DIABETIC MICE	55
6.1	Method	56
6.2	Experimental Methods	59
6.2.1	Accuracy Assessment	60
6.2.2	Reproducibility Study	61
6.3	Results	61
6.4	Discussion	65
7	A COMBINED MACHINE-LEARNING AND GRAPH-THEORETIC APPROACH FOR THE SEGMENTATION OF RETINAL SURFACES IN SD-OCT VOLUMES	69
7.1	Methods	70
7.1.1	Segmentation Process	71
7.2	Experimental Methods	74
7.2.1	Data Description	74
7.2.2	Accuracy Evaluation	76
7.3	Results	77
7.3.1	Data Obtained from Human Subjects	77
7.3.2	Data Obtained from Mice	79
7.3.3	Data Obtained from Canines	81
7.4	Discussion and Conclusion	82
8	INCORPORATION OF LEARNED SHAPE PRIORS INTO A GRAPH-THEORETIC APPROACH WITH APPLICATION TO THE 3-D SEGMENTATION OF INTRARETINAL SURFACES IN SD-OCT VOLUMES OF MICE	87
8.1	Method	87
8.2	Experimental Methods	90
8.3	Results	91
8.3.1	Optimization of Weighting Parameter a	91
8.3.2	Segmentation Accuracy	91
8.4	Discussion	94
9	A GRAPH-BASED APPROACH FOR THE SEGMENTATION OF MULTIPLE SURFACES WITH A SHARED HOLE	97
9.1	Method	99
9.2	Segmenting the Projected Hole Boundary Columns	102
9.3	Simultaneous Segmentation of the Surfaces and Their Shared Hole	102

9.3.1	Feasibility Constraints	103
9.3.2	Cost Function	106
9.4	Experimental Methods	109
9.5	Results	111
9.6	Conclusion & Discussion	120
10	INVESTIGATION OF TEXTURAL MARKERS DERIVED FROM SD-OCT IMAGES THAT MAY PREDICT STRUCTURAL AND FUNC- TIONAL CHANGE IN DEGENERATIVE DISEASES	122
10.1	Method	122
10.1.1	Human Data	123
10.1.2	Canine Data	124
10.2	Results	125
10.2.1	Human Data	125
10.2.2	Canine Data	126
10.3	Discussion & Conclusion	127
11	CONCLUSION	129
	REFERENCES	132

LIST OF TABLES

Table

3.1	Summary of previously proposed methods for the segmentation of intra-retinal layers in OCT images.	18
4.1	Mean unsigned difference seen in the original, partially corrected and final artifact-corrected images.	37
5.1	Average mean unsigned border positioning error [†] computed on 10 glaucomatous datasets.	51
6.1	Signed and unsigned border position error [†] (mean \pm standard deviation in μm) for each of the 10 segmented intra-retinal surfaces.	64
6.2	Summary of the reproducibility analysis for four retinal layers.	64
6.3	Summary of the reproducibility analysis for 11 retinal layers.	65
7.1	Unsigned border position error (mean \pm SD) in μm computed on the human dataset.	78
7.2	Unsigned border position error (mean \pm SD) in μm obtained on the mice data.	80
7.3	Overall unsigned border position error (mean \pm SD) in microns obtained on the 19 canine scans.	82
8.1	Unsigned border position error observed for the 10 retinal surfaces for the range of value of a tested.	93
8.2	Accuracy assessment of segmentation results when incorporating learned shape priors.	94
9.1	Summary of accuracy assessment of NCO segmentation	113
10.1	Mean thickness of RNFL and GCL in predefined sectors	126
10.2	Mean thickness of NF+GC+IPL in predefined sectors in canine data. . .	128

LIST OF FIGURES

Figure	
2.1	Retinal layers visible in SD-OCT images. 9
2.2	The neural canal opening as seen in SD-OCT images. 11
3.1	Varying feasibility constraints. 14
3.2	Illustration of the cost function utilized by a graph-theoretic approach. . 16
4.1	Examples of commonly acquired SD-OCT images. 25
4.2	Schematics showing acquisition of OCT datasets with orthogonal fast scanning axes from the same patient. 27
4.3	Selected slices from an OCT dataset showing the two surfaces segmented in the original volume. 28
4.4	Overview of method to determine the flattening plane. 30
4.5	Fundus photograph and its corresponding disparity maps. 34
4.6	Estimation of the shape of the eye from paired OCT datasets with orthogonal fast scanning axes. 36
4.7	Example of the surface used to create the reference plane. 38
4.8	Central B_f -Scan and B_s -scan slices from an OCT dataset before and after flattening. 39
4.9	An ONH-centered dataset before and after the axial artifact correction. . 40
5.1	Two example of diseased scans that proved problematic for the baseline approach that did not incorporate any learned features. 44
5.2	Illustration of the interacting surfaces seen in SD-OCT retinal images. . . 45
5.3	A schematic sketch of the proposed segmentation method that incorporated learned textural features into the cost function utilized by the graph-theoretic method. 48
5.4	Probability maps obtained for the various region with and without grouping on the basis of similar properties 50
5.5	Segmentation results. 52

5.6	Segmentation results.	53
6.1	Slice from an SD-OCT volume showing 10 retinal surfaces	55
6.2	Overview of segmentation process used to segment mice SD-OCT volumes.	57
6.3	Illustration of the segmentation process for the outer retinal surfaces	57
6.4	Illustration of the segmentation process for the inner retinal surfaces.	58
6.5	Projection images obtained from volumetric SD-OCT images of mice.	59
6.6	Labeled sectoral map within which the reproducibility of thickness measurements were assessed.	62
6.7	Segmentation results of the mice SD-OCT scans.	63
6.8	The thickness differences computed in four retinal layers.	66
6.9	The coefficient of variation (CV%) computed in four retinal layers.	67
7.1	Examples of images used to evaluate the efficacy of the proposed method.	70
7.2	Probability maps obtained for the three categories of regions and two categories of surfaces for scans acquired from a human, a mouse and a canine.	72
7.3	The evaluation criteria for scans of mice.	74
7.4	The segmentation accuracy evaluation criteria for scans of canines.	75
7.5	Segmentation results on the human dataset.	79
7.6	Segmentation results obtained for the mice dataset.	81
7.7	Segmentation results obtained on the dog dataset.	83
8.1	A central slice from an SD-OCT volume of a mouse showing the 10 surfaces segmented.	88
8.2	OOverview of the segmentation process used to segment 10 retinal surfaces in the mice SD-OCT scans.	89
8.3	Graph of error vs. tested values of a	92
8.4	Segmentation results obtained after incorporating learned shape priors.	95
9.1	Views of the neural canal opening in 3-D and 2-D.	98

9.2	Example of surface segmentation errors frequently seen at the NCO. . . .	99
9.3	Representation of three regions created by the ‘surfaces + hole’ configuration.	100
9.4	Schematic showing the segmentation process that includes a 2-D and 3-D component.	101
9.5	Result of stationary wavelet transform-based filtering.	103
9.6	Constraints used for regions \mathcal{B} and \mathcal{O}	104
9.7	Polar transformation of volumetric images.	107
9.8	Example of scan with substantial artifacts.	110
9.9	A segmentation result.	114
9.10	Another example of the NCO segmentation.	115
9.11	Correlation plots of MRW measurements obtained using manual delineations and the automated methods.	117
9.12	Bland-Altman plots of MRW measurements obtained using manual delineations and the automated methods.	118
9.13	Example of where the iterative method erroneously detected the end of the border tissue instead of the NCO.	119
10.1	Sector maps used to analyze human scans.	124
10.2	The sector map used to analyze the dog scans.	125
10.3	Bar plot of features obtained from normals, suspects and glaucoma patients.	127

CHAPTER 1 INTRODUCTION

Optical coherence tomography (OCT) since its introduction in 1991 [1], has found widespread use in the detection and management of ocular diseases. This noninvasive imaging modality is capable of acquiring sub-micrometer resolution images of the retina, whereby the different retinal structures can be quantitatively studied. The new availability of the spectral-domain (SD-OCT) scanners, which became commercially available in 2007, has further provided volumetric (3-D) images of the retina.

The automated segmentation of structures in these images is, however, far from trivial as the images are prone to speckle noise and often show large artifacts caused by motion and acquisition errors. The presence of disease can also cause low contrast in these images, making it difficult to identify structures of interest. Structural variation, inherent even in a normative set, can be significantly larger in the presence of disease or when the research study is extended to an animal model involving mice, canines or monkeys. Another challenging aspect of the problem is the large size of the volumetric images: one of the imaging protocols available on the Cirrus SD-OCT scanner (Carl Zeiss Meditec Inc.) acquires images from a region $6\text{mm} \times 6\text{mm} \times 2\text{mm}$ and is comprised of $200 \times 200 \times 1024$ voxels. Similarly, the Bioptogen animal (mice) scanners can acquire images from a region $1.4\text{mm} \times 1.4\text{mm} \times 2\text{mm}$ that contain $400 \times 400 \times 1024$ voxels (four times larger than Cirrus scans). The spatial context available in these 3-D scans can be of crucial importance to automated segmentation methods, as structures that are ambiguous in a single 2-D slice may be better defined in 3-D. However, in order to leverage this contextual information, the automated methods developed must be capable of handling the large size of the scans. Some commercial scanners are equipped to provide automated thickness measurements of a few retinal layers. For instance, the Cirrus scanner provides measurements and thickness maps of the retinal nerve fiber layer (RNFL), the ganglion cell layer (GCL) and

the total retinal thickness. These automated segmentation methods, can however, fail on occasion. Moreover, only a limited amount of this information is available for export off the scanner, such as the mean thickness of a layer within predefined regions. While these parameters have proven to be useful in a clinical setting, research studies often focus on more than three layers and having access to the complete layer segmentations would allow for further statistical analysis.

A number of the early approaches [2–10] proposed for the segmentation of intra-retinal layers in SD-OCT images were 2-D methods that relied heavily on edge profiles and were largely validated on normative datasets rather than on diseased data. The large variability in contrast in these images between subjects as well as structural variation introduced by disease has encouraged the use of machine learning-based methods, such as, pixel classification [11–13], shape models (parametric [14, 15] and non-parametric [16]) and active appearance models [17]. Graph-based methods such as dynamic programming [18] and shortest-path algorithms [19, 20] have also been described in the literature. All of these methods have been described in 2-D, and while the pixel classification, shape and appearance models and active contour methods can be extended to 3-D, the same cannot be said of the graph-based approaches that used shortest-path algorithms.

The first 3-D method described to address this segmentation problem was proposed by Garvin *et al.* [21]. This method employed a graph-theoretic approach [22, 23] for the simultaneous segmentation of the intra-retinal layers in SD-OCT images centered on the macula. The segmentation of the set of feasible surfaces was transformed into a minimum-closure problem within a vertex-weighted graph. The surface set feasibility is expressed using expected surface smoothness and the distance between surfaces, and is learned from a training set [24]. The surface smoothness is expressed in two axial directions, thereby incorporating 3-D contextual information, while the surface interaction constraints broadly incorporates shape information. The cost func-

tion, which was a weighted combination of a gradient and a “volumetric” region term, was used to assign weights to the nodes in the weighted graph, while the structure of the graph itself enforced the constraints. The problem is then transformed into a minimum $s - t$ cut problem, thereby allowing for the simultaneous segmentation of the surfaces while ensuring that the result is globally optimal with respect to the cost function. Although the method boasts a polynomial time complexity that depends on the size of the graph, the run-time can be substantial when segmenting multiple surfaces in the original volumes. Lee *et al.* [25] proposed the use of a multi-resolution graph-theoretic approach that specifically addressed this problem and provided a substantial reduction in run time. Another extension to the method was proposed by Song *et al.* [26] who described the use of learned shape priors that augmented the graph structure creating a greater influence on surface topology than seen in previous formulations of the method. This method showed improvements over previously published results on a set of SD-OCT scans obtained from normal subjects [21]; however, its impact on a diseased set has not been similarly compared. The graph-theoretic approach has also been proposed for the segmentation of intra-retinal layers at the optic nerve head [27] as well as the neural canal opening [28] in SD-OCT images.

The cost functions used in the aforementioned approaches have mostly been designed by hand and incorporated little in terms of learned features. However, the graph-theoretic approach has seen such an extension in a different non-ophthalmic application. Yin *et al.* [29] used the graph-theoretic approach for the segmentation of cartilage and bones in MR scans of the human knee-joint. Here, the problem was formulated to simultaneously segment multiple interacting objects and surfaces. The cost function incorporated learned Haar and geometric features learned from a training set. However, it is important to note that this method used a pre-segmentation step to find the approximate boundaries of the bones. Furthermore, the influence of the incorporated features were not statistically assessed.

Our objective in this work, is to build on these key concepts and establish a structured framework for the machine learning-based design of the various components of the graph-theoretic approach, namely the cost function (comprised of in-region and on-surface terms) and the structure of the graph. Additionally, the framework will also assess the impact of the incorporated learned components on the segmentation accuracy. Thus, given a training set, the framework will allow for the graph-theoretic approach to be tailored to a specific problem.

The following specific aims were used to attain these objectives:

Aim 1: Develop and evaluate a 3-D automated method for the correction of axial artifacts in SD-OCT volumes. The artifacts commonly seen in SD-OCT scans are known to differ along the different axes of the image. We incorporated this *a priori* information into the method in order to correct these artifacts and bring the datasets into a more consistent shape. This reversible process thereby allows for more accurate shape information to be learned from the datasets.

Aim 2: Develop a machine-learning based framework for the automated design of components of a graph-theoretic approach for the automated segmentation of intra-retinal surfaces from SD-OCT images. Textural and shape features were learned from a training set and incorporated into the cost function and the graph structure, respectively, for the segmentation of intra-retinal layers in SD-OCT images obtained from human and animal subjects. The inclusion of the learned components into the graph-theoretic method was conducted in a sequential manner and its impact on the segmentation results gauged. The final optimal segmentation framework, therefore, consists of the learned components that significantly improve the accuracy of the method.

Aim 3: Develop a formulation of the graph-theoretic approach for the si-

multaneous segmentation of multiple surfaces with a shared hole. The “hole” is defined as a 3-D planar disruption in the surfaces. Therefore, the cost functions and the graph structure used in this formulation were adapted to meet these new requirements while continuing to ensure the global optimality of the surfaces and the shared hole. A machine learning-based approach was used to design the cost functions appropriately. The method was applied to the segmentation of retinal surfaces and the neural canal opening in SD-OCT images.

Aim 4: Investigate the hypothesis that textural markers derived from SD-OCT images are quantifiable before the occurrence of structural and functional change seen in degenerative diseases that affect retinal ganglion cells. The RNFL and the GCL are structures that are known to be affected by degenerative diseases such as glaucoma and diabetic retinopathy. Therefore, textural features within these layers were statistically studied in SD-OCT images obtained from normal and subjects with early glaucoma to find possible textural precursors to the disease. The study was also conducted using longitudinal data obtained from canines.

The segmentation of retinal structures in SD-OCT images is the focus in this work, however the methods (specifically Aims 2 and 3) have been formulated for wider applicability.

1.1 Thesis Overview

This thesis is divided up into 11 chapters, where:

- Chapter 2 discusses the clinical motivations behind this research as well as background information regarding the structures that are segmented and studied in this work.

- Chapter 3 provides an overview of the graph-theoretic approach [21, 23, 30] utilized in this work as well as prior approaches proposed for the automated segmentation of structures in SD-OCT images.
- Chapter 4 summarizes the methods and validation completed for the correction of axial artifacts in SD-OCT images.
- Chapters 5 through 8 summarize the methods and validation completed for the incorporation of learned components into the graph-theoretic method. Chapter 5 details the incorporation of learned textural features into the cost function, applied to the segmentation of retinal surfaces in scans obtained from glaucoma patients. Chapter 6 summarizes a reproducibility study conducted on the segmentation of retinal surfaces in mice scans using a baseline approach that did not incorporate any learned components. Chapter 7 describes the methods and validation completed for the automated design of the cost function utilized by the graph-theoretic approach, applied to the segmentation of retinal surfaces in human, mice and canine scans. Chapter 8 summarizes the methods and validation completed for the incorporation of learned shape priors into the graph-theoretic approach applied to the segmentation of mice scans.
- Chapter 9 summarizes the methods and validation completed for the simultaneous segmentation of surfaces and their shared hole, applied to the segmentation of retinal surfaces and the neural canal opening in human SD-OCT scans.
- Chapter 10 summarizes the results of the investigation of the hypothesis that textural markers can be used to predict structural and functional change associated with degenerative diseases that affect retinal ganglion cells.
- Concluding remarks and possible of future avenues of research are briefly discussed in Chapter 11.

CHAPTER 2 CLINICAL BACKGROUND

As a significant portion of the diseased scans use to evaluate the methods came from glaucoma studies, the clinical aspects of this disease as well as its detection and management are discussed briefly, followed by a description of relevant retinal structures.

2.1 Clinical Motivation

Glaucoma is the second leading cause of blindness in the United States and affects 1-2% of the U.S. population [31,32]. The disease is often characterized by an increase in intraocular pressure (IOP) and in the past, was often defined by this symptom. But recent studies have shown that up to a sixth of patients with glaucomatous damage never show increased IOP values, and one-third to one-half of patients do not have elevated IOP values in the initial stages of the disease. Glaucoma is now defined as a progressive disease of the optic nerve that can cause severe vision loss if left untreated. As this disease affects peripheral vision, it often goes unnoticed by the patient until considerable damage has already occurred. The disease is treatable and permanent vision loss can be avoided if the disease is detected early. Clinically, visual field tests and planimetry based on stereo fundus photographs are used for the diagnosis of glaucoma. However, these tests are subjective and have low reproducibility, increasing the need for automated, quantifiable and reproducible methods that can be used for the early diagnosis and management of the disease.

Structurally the disease is known to affect the retinal nerve fiber layer (RNFL) and the ganglion cell layer (GCL), which is seen to reduce in thickness as the disease progresses. This thinning can be quantitatively measured in optical coherence tomography (OCT) images, and thus, this imaging modality has begun to find widespread use in the diagnosis and management of the disease. Quantitative comparisons between the RNFL thickness measured in OCT images acquired from normal and glau-

comatous eyes have been shown to have significant differences, which can be used to detect the presence of the disease [6, 33–35] and has also found use in automated screenings [36]. Further investigations by Garvin *et al.* [37] and Lee *et al.* [38] have demonstrated a strong correlation between thicknesses of retinal layers along fiber bundle tracts and thickness of the neural rim and the peripapillary nerve fiber layer, thus, further emphasizing the potential use of OCT as a quantitative marker of disease progression.

The thinning of the RNFL as well as the glaucomatous cupping seen in stereo fundus photographs was thought to be caused solely by the axonal damage to the retinal ganglion cells. However, it has since been proposed that the glaucomatous cupping is caused by a combination of axonal cell loss and damage to underlying connective tissue caused by increased intraocular pressure (IOP) [39, 40]. This raised interest in the the neural canal opening (NCO), a three-dimensional structure discernible in SD-OCT images, as it is unaffected by glaucomatous cupping. The NCO, can therefore, be used as a stable reference plane [41, 42] from which optic nerve head parameters can be computed for the detection and tracking of glaucomatous change [43].

OCT imaging allows for the quantitative assessments of critical structures in the retina, which has lead to its widespread use in the diagnosis and management of a number of ocular diseases. Automated analysis techniques would not only augment current disease management methods but would also allow for automated screenings that could detect the diseases early enough to prevent vision loss.

2.2 Retinal Structures

The retina is a complex structure consisting of several layers of neurons, and forms a continuous lining inside the eye and is the structure responsible for vision. The light sensitive photoreceptors, known as rods and cones, convert light into electrical signals that are relayed to the brain via the optic nerve. The rods and cones face the incoming light, but only sense the light that is reflected back towards them. When stimulated

by light, these cells spark a chemical reaction that is propagated through the layers until it finally reaches the ganglion cells, and thereafter transmitted to the brain. The macular region contains the largest number of photoreceptors, while the optic nerve head is the point from which the optic nerve leaves the eye.

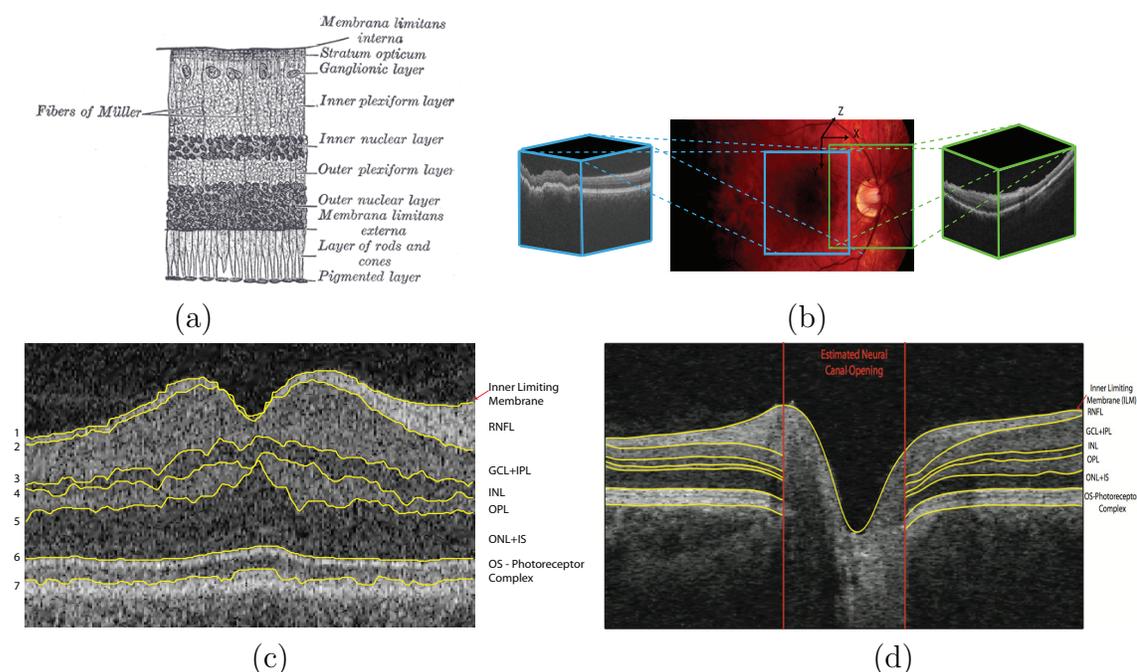


Figure 2.1: (a) A lithograph plate from Gray's Anatomy (originally published in 1918 and has since lapsed into public domain) showing the cross-sectional view of the ten retinal ten layers. (b) Fundus photograph depicting the regions scanned in macula-centered and ONH-centered OCT images. (c) A central xz -slice from a macular SD-OCT dataset showing some of the commonly segmented surfaces. (d) A central xz -slice from an ONH-centered dataset showing the various layers of the retina as well as the neural canal opening.

The structures and layers seen in the retina are as illustrated in Fig. 2.1, and are described briefly below:

- Inner Limiting Membrane (ILM) is the boundary between the retina and the vitreous body.

- Retinal Nerve Fiber Layer (RNFL) consists of axons of the ganglion cells that form the optic nerve.
- Ganglion Cell Layer (GCL) contains the nuclei of the ganglion cells.
- Inner Plexiform Layer (IPL) contains axons of bipolar and amacrine cells, and the dendrites of the ganglion cells.
- Inner Nuclear Layer (INL) contains the nuclei of horizontal, bipolar, amacrine and Müller cells.
- Outer Plexiform Layer (OPL) consists of photoreceptor axons, and dendrites of horizontal and bipolar cells.
- Outer Nuclear Layer (ONL) consists of the cells bodies of the photoreceptor (both rods and cones) cells.
- The inner segments (IS) of the photoreceptors appears as a dark band in OCT images while the junction between the inner and outer segments (IS/OS junction) appears as a bright surface.
- Retinal Pigment Epithelium (RPE) is a single layer of cells between the retina and the choroid.

In OCT images, the surface between the GCL and IPL layers is not consistently seen and is therefore, often not segmented. Fig. 2.1(c) shows a cross-sectional slice from an SD-OCT image centered on the macula, which is the point with the highest density of cones and is associated with high visual acuity. The optic nerve head (Fig. 2.1(d)) is the location where the ganglion cell axons leave the eye to form the optic nerve. There are no photosensitive cells at this region and is thus, insensitive to light and is called the *blind spot*. The optic nerve head also shows a number of large blood

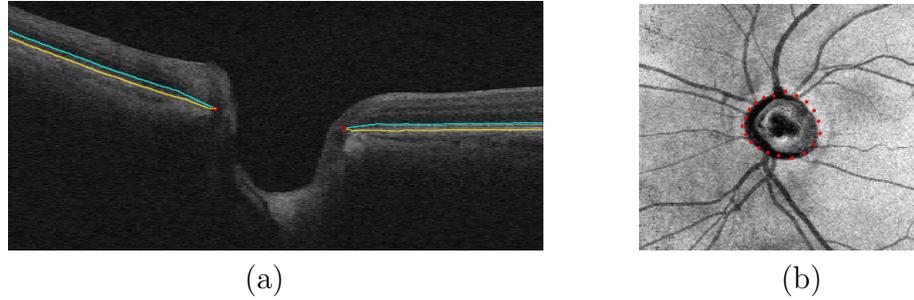


Figure 2.2: (a) A central slice from an SD-OCT scan from a nonhuman primate eye showing two surfaces in the RPE-complex and the NCO. (b) A projection image created from the SD-OCT volume showing the NCO as manually delineated by an independent trained expert.

vessels as this is the location from which the vessels that supply the ocular tissue enter and leave the eye.

The neural canal opening (NCO) (see Fig. 2.2), is a three-dimensional planar structure [44] that occurs at the level of the retinal pigment epithelium (RPE) and the Bruch's membrane and marks the anatomic entrance to the neural canal. This structure is referred to as the Bruch's membrane opening (BMO) in the literature and in non-human primate eyes, the NCO colocalizes to the BMO. The Bruch's membrane is the anterior surface of the choroid and the BMO is the location at which the optic nerve passes through this membrane. The anatomy in humans is more complicated than in primates, and there is some ambiguity about where these separate structures can truly be discerned in SD-OCT images obtained from human subjects. This structure shows little change in the presence of early glaucomatous [40,41] and it has been used as a stable reference plane from which measurements were made in histomorphometric reconstructions of primate eyes. These studies showed promising results, encouraging the use of the same for the detection and tracking of glaucoma in human subjects.

CHAPTER 3 TECHNICAL BACKGROUND

The graph-theoretic approach is utilized extensively in this work and is therefore, outlined below in Section 3.1. The previous contributions in the area of the segmentation of ophthalmic structures in SD-OCT images are discussed in Section 3.2 to further motivate the chosen strategy.

3.1 Graph-Theoretic Method

The graph-theoretic approach proposed by Li *et al.* [23] allowed for the simultaneous segmentation of multiple interacting surfaces, while ensuring a globally optimal solution. The method converts the segmentation problem into an optimization one where the aim is to find a feasible set of surfaces (with the minimum cost) that form a *closed set* (a set of nodes in a graph where no edges leave the set) in a graph. The minimum-cost closed set problem is in turn converted into a minimum $s - t$ cut problem. The graph is constructed using two main descriptors (described below) - the constraints that describe the surface set feasibility and the cost function with respect to which the minimum-cost closed set is computed.

3.1.1 Feasibility Constraints

Consider a volumetric image $\mathcal{I}(x, y, z)$ with dimensions $X \times Y \times Z$, where the surface \mathcal{S} is a function $\mathcal{S}(x, y) \in \mathbf{z} = \{0, \dots, Z-1\}$, where $x \in \mathbf{x} = \{0, \dots, X-1\}$ and $y \in \mathbf{y} = \{0, \dots, Y-1\}$, implying that the surface intersects each (x, y) column in a single location exactly once. The feasibility of the set of surfaces is represented using smoothness and surface-interaction constraints. In the original formulation [22, 23], the smoothness constraints Δ_x and Δ_y represented the maximum permitted distance between two adjacent voxels in the x and y directions, respectively. For instance, if $\mathcal{I}(x, y, z_1)$ and $\mathcal{I}(x + 1, y, z_2)$ are two adjacent voxels on a surface, then $|z_1 - z_2| \leq \Delta_x$. Δ_y is expressed in a similar manner for adjacent voxels in the y -

direction. Haeker/Garvin *et al.* [21, 45] proposed using smoothness constraints that varied with respect to the location of each voxel. Thus, given two adjacent columns p and q within a predefined neighborhood \mathcal{N} (such as a 4-neighbor relationship), the smoothness constraint can be defined as follows:

$$-\Delta_{p,q}^u \leq f(p) - f(q) \leq \Delta_{p,q}^l, \quad (3.1)$$

where, $\Delta_{p,q}^u$ and $\Delta_{p,q}^l$ represent the maximum allowed change in z -value in the upward and downward direction, respectively.

The surface-interaction constraints define the relationship between the surfaces, such as the ordering of the surfaces as well as the expected distances between them. For instance, it may be known that surface i is above surface j and that the minimum and maximum allowed distances between them is $\delta_{i,j}^{min}$ and $\delta_{i,j}^{max}$, respectively. The surface-interaction constraints, like the smoothness constraints, can be defined as constant values for pairs of adjacent surfaces [23] or they can be allowed to vary as a function of the column location, [21, 45] and can be expressed as follows:

$$\delta_{i,j}^{min}(p) \leq \mathcal{S}_j(p) - \mathcal{S}_i(p) \leq \delta_{i,j}^{max}(p), \quad (3.2)$$

where, $\delta_{i,j}^{min}(p)$ and $\delta_{i,j}^{max}(p)$ represent the minimum and maximum allowed distance between surfaces i and j at location p . The location-based specification of the constraints allow for the incorporation of smoothness and surface-interaction constraints learned from a training set. Figs. 3.1(a) and (b) show the graph structure obtained when using varying smoothness and surface-interaction constraints.

In order to further influence surface topology, Song *et al.* [26, 46] proposed the use of additional weighted arcs as shape penalties. These were introduced in the form of additional arc weights with rational weights, unlike in the original formulation where

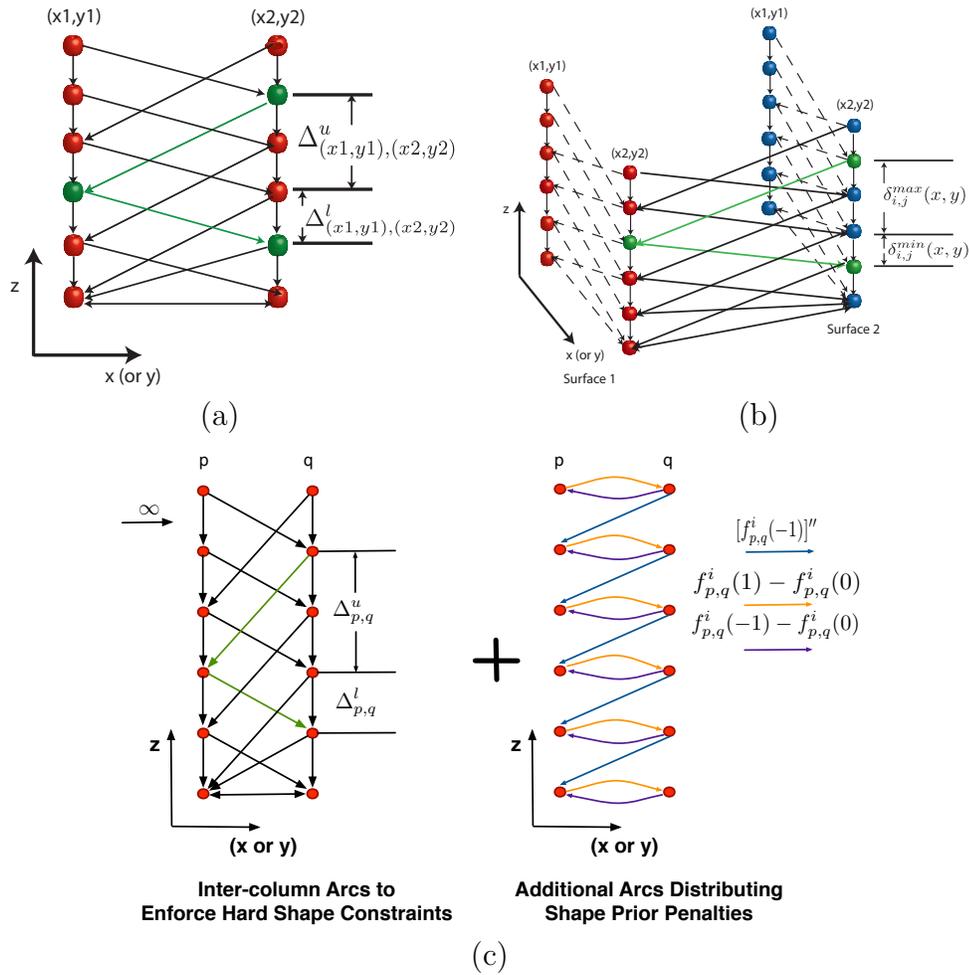


Figure 3.1: Feasibility constraints. Illustration of (a) varying smoothness constraints, (b) varying surface-interaction constraints and (c) incorporation of shape priors. Here, it is assumed that $m(p, q) = 0$, thus backward and forward arcs are added between (p, z) and (q, z) .

the arcs have infinite weights. The shape prior penalty is assumed to be a convex function $f((\mathcal{S}(p) - \mathcal{S}(q)) - m(p, q))$, where the deviation of the surface \mathcal{S} away from the mean shape $m(p, q)$ is penalized. Intuitively, this means that shape penalties are only imposed when $\mathcal{S}(p) - \mathcal{S}(q) \neq m(p, q)$.

The shape prior penalty is distributed between the columns of the graph based on the discrete second derivative of the function given by $f''(h) = [f(h+1) - f(h)] - [f(h) - f(h-1)]$. Since the function is assumed to be convex, $f''(h) \geq 0$. A shape penalty arc is added for each $h = (\mathcal{S}(p) - \mathcal{S}(q)) - m(p, q)$, where $-\Delta_{p,q}^u \leq (\mathcal{S}(p) - \mathcal{S}(q)) \leq \Delta_{p,q}^l$. If $f'(h) \geq 0$, nodes (p, z) and $(q, z - m(p, q) - h)$ are connected by an arc with weight of $f''(h)$. Similarly, if $f'(h) \leq 0$, nodes (p, z) and $(q, z + m(p, q) + h)$ are connected by an arc with weight of $f''(h)$. This construction ensures that the total weights of the arcs spanning between two neighboring columns p and q equals the shape prior penalty $f((\mathcal{S}(p) - \mathcal{S}(q)) - m(p, q))$. Fig. 3.1(c) shows the additional arcs that are added to enforce the shape penalty when $m(p, q) = 0$. In this situation, at (p, z) and (q, z) , $f'(0) = 0$, the forward arc $f''(0)$ is represented by $f''(0^+) = f(1) - f(0)$ and the backward arc $f''(0^-) = f(0) - f(-1)$. Thus, when $\mathcal{S}(p) - \mathcal{S}(q) = 0$ no penalty arcs are cut and the shape penalty is zero.

3.1.2 Cost Functions

The minimum cost set of feasible surfaces is computed with respect to a cost function, which can be designed using a combination of on-surface and region-based costs. In the original formulation [23], the in-region costs were derived from a small region around each voxel. [21,24] described a cost function that incorporates “true” regional information in addition to edge gradients. In such a scenario, n non-intersecting surfaces would divide the volume into $n + 1$ regions, thus associating each voxel with n on-surface costs and $n + 1$ in-region cost terms, as shown in Fig. 3.2. The on-surface cost terms reflect the unlikelihood that the voxel lies on surface i , $i \in \mathbf{i} = \{1, \dots, n\}$, while the in-region cost terms reflect the unlikelihood that the voxel belongs in region

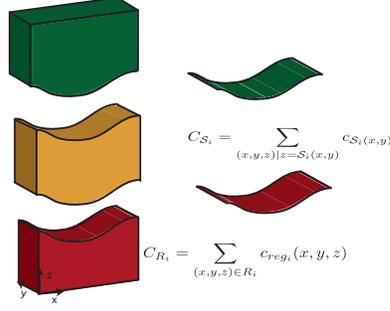


Figure 3.2: Illustration of the cost function, depicting the on-surface and in-region cost terms.

$i, i \in \mathbf{i} = \{0, \dots, n\}$. The total cost C_T , associated with the set of n surfaces can now be expressed as a sum of the on-surface and in-region cost terms:

$$C_T = \alpha \sum_{i=1}^n \sum_{(x,y,z)|z=S_i(x,y)} c_{surf_i}(x,y,z) + (1 - \alpha) \sum_{i=0}^n \sum_{(x,y,z) \in R_i} c_{reg_i}(x,y,z). \quad (3.3)$$

where, $c_{surf_i}(x,y,z)$ and $c_{reg_i}(x,y,z)$ represents the on-surface costs and in-region costs associated with the surface \mathcal{S}_i and region i , respectively and α is a weighting parameter.

When the soft shape constraints are incorporated into the method, the cost function now includes a shape penalty term C_{shape} . The shape prior penalties are defined using the second derivative of the penalty function, thereby requiring the function to be convex. Thus, a quadratic function could be utilized giving us the following representation of C_{shape} :

$$C_{shape} = \sum_{i=1}^n a_i \sum_{p,q \in \mathcal{N}} ((\mathcal{S}_i(p) - \mathcal{S}_i(q)) - m_i(p,q))^2 \quad (3.4)$$

where, $m_i(p,q)$ is the mean expected change between columns p and q for surface \mathcal{S}_i

and a_i is the weighting co-efficient.

3.1.3 Graph Construction

Sub-graphs can now be created for each of the surfaces, where the smoothness constraints and the shape prior penalties determine the structure of the graph and the vertex weights are determined by the in-region and on-surface cost terms. The sub-graphs are then connected using arcs that reflect the pre-defined surface-interaction constraints [21, 23, 45]. Multiple surfaces can now be detected simultaneously on this vertex-weighted graph by computing the minimum-cost closed set, through a minimum-cost $s - t$ cut [22, 23].

3.2 Review of Prior Approaches Proposed for the Segmentation of Retinal Surfaces in SD-OCT Images

The automated segmentation of intra-retinal layers in OCT images is a challenging problem and has been approached in numerous ways (summarized in Table 3.1). Some of the earliest methods were 2-D methods that relied heavily on edge information in the A-scan lines and used 2-D context to verify the segmentation [2–10]. These methods follow four main steps. 1) The datasets are filtered, 2) edges are detected in each A-scan line, 3) the large gradient that marks the location of the surface between the inner and outer segments of the photoreceptors is used to align the A-scan lines (flatten the dataset) and 4) the detected edges are verified using 2-D contextual information and then connected to form a smooth surface. Additional features to this schema included the mapping of grayscale intensities [7], the averaging of multiple A-scans [9,10], the use of an iterative method that refined the result [8] and the use of a Markov model to correct errors in the peak detection [47]. As gradient information in OCT images can be unreliable, the use of iterative statistical regression and polynomial line fitting [48], as well as Canny edge detection on filtered images [49] have also been proposed to further improve the performance of the method. These

Table 3.1: Summary of previously proposed methods for the segmentation of intraretinal layers in OCT images.

Type	Method	Citations
1-D	Edge profile analysis	Koozekanani (2001), Fernandez (2005), Shahidi (2005), Ishikawa (2005), Chan (2006), Baroni (2007), Bagci (2008), Tan (2008, 2009), Fabritius (2009), Zhang (2012)
	Active contours	Mujat (2005), Mishra (2009), Yazdanpanah (2011)
2-D	Machine Learning - pixel classification	Zawadzki (2007), Rossant (2009), Vermeer (2011), Lang (2013)
	Machine learning - shape/texture models	Kajic (2010), Rathke (2011)
	Graph-based shortest path algorithms	Chiu (2010, 2012), Yang (2010)
3-D	Graph-theoretic approach	Garvin (2009), Lee (2009), Antony (2010), Song (2010), Dufour (2013), Lang (2013)

methods, however, do not incorporate much in the way of contextual information and the methods' ability to deal with diseased datasets where the images sometimes show poor contrast have not been adequately demonstrated.

Machine-learning based approaches have also been proposed [11–13, 50], where learned features from the retinal layers are classified, thus, segmenting the images. Rossant *et al.* [12] used such an approach, where k-means clustering was used to segment the inner retinal layers. The ILM and IS/OS surfaces were however, segmented using peak detection in edge profiles. Furthermore, the method assumed that the image is a 2-D slice centered on the macula and used the location of the fovea to specifically correct errors that occur at the fovea. Zawadzki *et al.* [11] proposed training a support vector machine (SVM) to find regions of interest in SD-OCT images and demonstrated the method on a variety of normal and diseased datasets. It is important to note however, that the method was not completely automated and required a user to train the SVM using an interface. Vermeer *et al.* [13], also proposed the

use of an SVM classifier trained on Haar-like features obtained from manually traced slices. As the classification result can be pixelated, the final segmentation is obtained after a post-processing step using level sets. The level set function incorporates gradients and surface smoothness constraints in addition to the labels obtained from the SVM, and was a 3-D function; however, the features themselves were generated from single A-scan lines. Lang *et al.* [50] proposed a method that utilized a random forest to learn the properties of the surfaces, where the final segmentation of the surfaces was obtained sequentially by finding the ordered locations along each A-scan that had the highest probability of being a surface. This method used the relative distance from a known structure as a feature in order to detect the surfaces correctly, but this increases the dependence of the method on the training set. Moreover, this method detected the surfaces in a sequential manner, which in turn increases the risk of propagating errors introduced into the segmentation at earlier steps.

Other approaches to the problem have employed active contours [14, 15, 51] in an effort to incorporate shape priors and regional information in addition to edge gradients. Mujat *et al.* [51] found the RNFL using a filtered image and a gradient image to minimize an energy function. While this makes the method robust to speckle noise, it remains dependent on gradient information. Mishra *et al.* [14] utilized a two-step approach, that initialized the location of the layers using dynamic programming, which was then followed by the fitting of an energy minimizing spline. However, the method's strength and weakness lies in its use of dynamic programming, which is used to simultaneously find an initial location of the surfaces (1-D lines) in the 2-D scans. Extending the same to find to surfaces in a 3-D volume is difficult and computationally extremely expensive. Yazdanpanah *et al.* [15], however, relied on region information rather than edge information and used a shape prior to reduce segmentation errors. The shape prior is a circle, and the structures are assumed to have a smooth arc-like shape. The use of a parametric shape prior may be adequate in normal scans, but

could prove to be far too limiting in the presence of disease-induced change.

Other model based approaches have used non-parametric shape models. Kajic *et al.* [17] proposed the use of an active appearance model (AAM) for the segmentation of retinal surfaces. The method used a large number of spatial and textural features learned from a large training set to create the AAM. However, the method is limited by the training set, which the authors were able to obtain using using a large number of manual tracers. A similar approach was proposed by Rathke *et al.* [16] who proposed the use of a probabilistic principal component analysis (PPCA) model to estimate the shape variations along each column as well as the overall global shape. The method also incorporated a local appearance model. The column-wise shape model was intended to allow the parallel segmentation of the surfaces, but showed poorer results than the global shape model. This method however, was designed for 2-D circular scans, where the shape variations (column-wise and globally) were limited and the extension to 3-D is far from trivial.

Graph-based approaches, such as shortest path algorithms and dynamic programming methods have also been employed for the segmentation of intra-retinal layers. Chiu *et al.* [19,20] described a method that utilized Dijkstra's shortest path algorithm to segment retinal surfaces. The beginning and end points of each surface was found in an initialization step, followed by the detection of the "shortest path", that provided the minimum-cost surface in the volume. The method is essentially a 2-D method, and the segmentations in the 2-D scans are stitched together to produce a surface. Yang *et al.* [18] described a similar approach that used a dynamic programming shortest path algorithm. Canny edge detectors with thresholds tailored to each of the surfaces of interest are used to generate gradient maps for the surfaces. Graphs are created for each of the surfaces of interest, where the nodes in the graphs are assigned values from the gradient maps. The "shortest path" obtained using a shortest-path algorithm was then used to obtain the surfaces. The Canny edge detector is known to

be robust to noise and the hysteresis loop does help reduce the dependence on edge gradients. However, the dependence on gradient information was not eliminated and no other features such as region texture or shape is used in this method. Also, this method was inherently 2-D and did not incorporate the 3-D contextual information available in volumetric scans.

Garvin *et al.* [21] presented a graph-theoretic approach [22, 23] (described in Section 3.1) that not only simultaneously segments intra-retinal surfaces in macular volumetric scans, but also ensures the global optimality of the solution. The method did not rely on shortest-path algorithms or dynamic programming, but formulates the problem as a minimum closure problem (solved using a minimum $s - t$ cut) that has polynomial time complexity. As described earlier in Section 3.1, the feasibility of the surfaces is represented in the smoothness and surface-interaction constraints and the optimization is done with respect to a cost function. In that work, Garvin *et al.* used constraints that were learned from a training set and the cost function consisted of on-surface and “true” volumetric regional cost terms. The on-surface cost terms were derived from directed gradients in the image and the regional costs were computed using fuzzy membership (Gaussian) functions that assigned each voxel with a probability of belonging to a particular region. The key advantage of this method over other methods described in the literature is the ability to segment multiple surfaces simultaneously in n -D. The graph structure created using the smoothness constraints inherently incorporates 3-D contextual information, while the surface-interaction constraints provides a broad “shape” context with respect to adjacent surfaces. Moreover, the final set of segmented surfaces is assured to be globally optimal with respect to the cost function.

This method has since been proposed for the segmentation of intra-retinal layers in ONH-centered scans [25, 27], as well as structures in fundus photographs [52]. The multi-resolution approach proposed by Lee *et al.* [25] in particular, provided a

significant saving in run time. The cost functions used in these formulations were designed by hand and incorporated little in the form of learned features. However, the application of this graph-theoretic approach is not limited to ophthalmology and has been proposed for the segmentation of structures in CT [46, 53] and MR [29, 54] where learned features such as texture and shape priors have been incorporated into the method. Yin *et al.* [29] formulated the graph-theoretic approach for the segmentation of multiple objects and multiple surfaces in MR images of the human knee joint whereby three bones and the interacting cartilage can be segmented. In that work, textural (Haar) and geometric features were incorporated into the cost function. The method also used a pre-segmentation step to find the approximate locations of the bones, which is necessary to construct the graph. Lang *et al.* [55] proposed a similar method for the segmentation of retinal surfaces in human SD-OCT scans, where the surface properties were learned from a training set. A random forest [56] was then used to create probability maps that were incorporated into the cost function. However, this method did not design or utilize an in-region component in the cost function. Song *et al.* [30] described the use of learned shape priors that provided greater influence over the surface topology. Here, additional weighted arcs were introduced into the graph structure in order to penalize deviations from the learned shape priors. Dufour *et al.* [57] also proposed a method that utilized hard and soft shape constraints to simultaneously segment multiple surfaces in retinal SD-OCT images. The graph described in this method utilized a larger neighborhood definition and therefore, incorporated shape priors that were not limited to immediately adjacent column locations in the graph. This amounted to downsampling the image and the graph built in this way, was therefore, smaller and had a smaller memory requirement.

Thus, we see that the graph-theoretic approach is adaptable to new geometric configurations and easily incorporates learned features (textural and shape), however the extensions described thus far have not statistically assessed the impact of the learned

component on the segmentation accuracy. Furthermore, when segmenting multiple surfaces as in the case of retinal surfaces, the cost function design can be challenging due to the presence of similar regions. The framework proposed in this work will address these aspects by 1) detailing a systematic process for the incorporation of learned features into the cost function and graph structure, 2) address the possible presence of similar regions and surface interfaces and finally 3) gauge the impact of the incorporated learned component on the system, whereby the graph-theoretic method can be tailored to specific applications.

CHAPTER 4

A 3-D AUTOMATED METHOD FOR THE CORRECTION OF AXIAL ARTIFACTS IN SD-OCT VOLUMES

The new SD-OCT [1, 58] scanners show a higher signal-to-noise ratio than the previous generation time-domain OCT scanners, while providing close-to-isotropic volumetric images (see Fig. 4.1(a)) of the retina. These images are, however, prone to large artifacts, which are thought to be caused by a number of factors such as motion of the eye and positioning of the camera. The sequential acquisition of the individual A-scans and the speed of the acquisition along the different axes, resulting in large artifactual shifts that differ in characteristic ways along the slow (B_s -scans or yz slices) and fast scanning axes (B_f -scans or xz slices). Fig. 4.1(b) and (c) show examples of B_s -scans and B_f -scans, respectively. The B_f -scans are less prone to motion artifacts as the speed of acquisition is high but often show large tilts, which are thought to be caused by imaging paraxially to the optical axis of the eye. The B_s -scans, however, are affected by motion (of the head and eye) and show characteristic high frequency artifacts as depicted in Fig. 4.1(c).

These artifacts not only make it difficult to visualize the data, but they also affect the further processing of the images. Of course, the retina in these scans is also curved due to the natural scleral curvature and while this is not an artifact, it can be advantageous to at least temporarily remove this curvature for some applications. For instance, segmentation algorithms that incorporate learned information about the layers ([17, 21]) do so by modeling the behavior of the surfaces of interest. But this is hard to do in the presence of unpredictable artifacts such as those that are seen in the B_s -scans of the OCT images. Surface behavior is also of interest clinically, where it could be used to compare pathological changes to normal data. Bringing the data into a consistent format is also important in 3-D registration applications [59], such as registering ONH to macular scans. Here, a predictable consistent shape can have a significant impact on the result. Thus, the need to correct these artifacts and bring

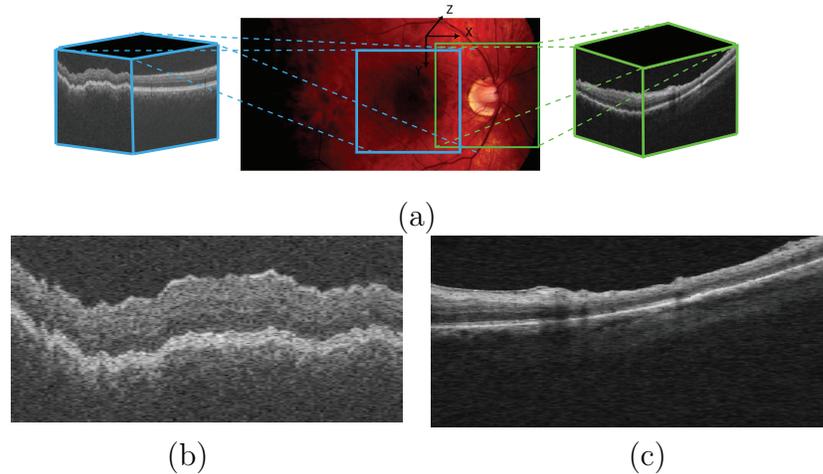


Figure 4.1: Examples of commonly acquired SD-OCT images. (a) Fundus photograph depicting the regions scanned in macula-centered and ONH-centered OCT images. (b) A typical central B_s -Scan from a macula-centered image showing artifacts characteristic to yz -slices. (c) A typical central B_f -scan (xz -slice) from an unprocessed ONH-centered volume, showing the tilt artifact commonly seen in these slices.

the dataset into a more consistent, predictable shape is compelling.

Since the correction process is intended to be a preprocessing step, it should alter as little as possible of the actual A-scans to avoid losing any information. The previously reported 1-D approaches [6] rely heavily on correlation between adjacent A-scans, leading to large errors in regions where the A-scans show rapid variability or low signal, such as vessel shadows. Numerous 2-D registration approaches [60–65] have also been proposed to tackle this problem. While these methods use more contextual information than the 1-D approaches, they do not incorporate 3-D contextual information. Furthermore, the registered images could be interpolated (unless the transform is purely translational), leading to altered image information. The new availability of orthogonal scans has also led to methods [66, 67] that incorporate information from both scans to realign the dataset and remove motion artifacts. This is an effective method that not only removes motion artifacts but also reconstructs the “true” shape of the retina. However, the application of these methods is restricted by

the availability of the orthogonal scans, which are not typically acquired clinically.

Garvin *et al.* [21], as a preprocessing step for a 3-D intraretinal segmentation algorithm, described a 3-D segmentation-based method that corrects motion artifacts by re-aligning the columns of the image with respect to a smooth “reference” plane. This reference plane is constructed by fitting a smoothing thin-plate spline (TPS) to a surface segmented in a lower resolution. The small number of control points and the large regularization term used in the spline-fit process reduces the dependence on the segmentation result, but the spline is not able to model the fast variations seen along the slow scanning axis. A smaller regularization term would have provided a closer fit to the control points, but this would increase the dependence of the artifact correction on the segmentation result.

The method described in this chapter is a segmentation-based method that address the characteristic artifacts along each axis separately while retaining the overall 3-D context. This is done by incorporating *a priori* information regarding the different artifacts seen along these two axial directions and correcting them using dual-stage thin-plate splines fitted to a segmented surface. Additionally, we also present a method to reconstruct the “true” scleral curvature (which is removed by the artifact correction method) given the new availability of orthogonal scans.

4.1 Data

The method was evaluated using two datasets containing macular-centered and ONH-centered OCT volumes:

1. Pairs of macula OCT scans were acquired from the same eye using orthogonal fast scanning axes on a SD-OCT1000 spectral-domain scanner (Topcon Corp., Tokyo, Japan). The scans were obtained from 9 normal subjects participating in the Rotterdam Study, which is a prospective population-based cohort study investigating age-related disorders. The study population consisted of

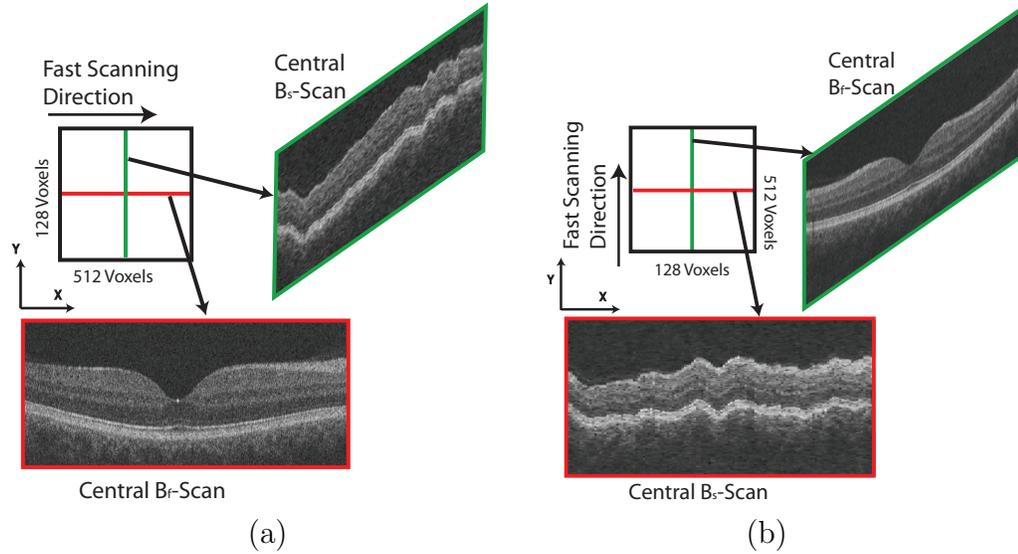


Figure 4.2: Schematics showing acquisition of OCT datasets with orthogonal fast scanning axes from the same patient. (a) Central B_f -scan and B_s -scan from an OCT image with horizontal fast-scanning axis. (b) Central B_f -scan and B_s -scan from an OCT image acquired from the same eye with vertical fast-scanning axis. Note that the B_f -scan from the first dataset comes from the same location as the B_s -scan from the second dataset.

7983 individuals aged 55 years and older living in the Ommoord district of Rotterdam, the Netherlands [68–70]. These OCT images, which were obtained as part of latest follow up in addition to other ophthalmic tests, had dimensions of $512 \times 128 \times 480$ voxels obtained from a $6 \times 6 \times 2$ mm³ region centered on the macula. Fig. 4.2 shows an example of a pair of macular-centered scans, where the central B_f -scan in Fig. 4.2(a) corresponds to the central B_s -scan in Fig. 4.2(b).

2. ONH-centered scans and stereo fundus photographs were obtained on the same day from 15 subjects (at the Glaucoma Clinic at the University of Iowa) that presented with glaucoma. The OCT scans were obtained from a Cirrus spectral-domain scanner (Carl Zeiss Meditec, Dublin, CA, USA), and had dimensions of $200 \times 200 \times 1024$ voxels obtained from a $6 \times 6 \times 2$ mm³ region centered on

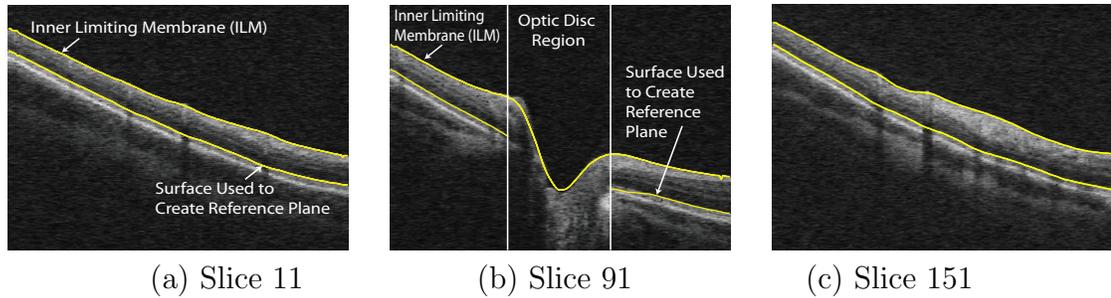


Figure 4.3: Selected slices from an OCT dataset showing the two surfaces segmented in the original volume. The reference plane is created by fitting a spline in two stages to the lower surface. The inner limiting membrane (ILM) as well as the lower surface are used in validation processes.

the ONH.

4.2 Methods

The artifact correction process needs a stable surface to which a thin-plate spline can be fit. We therefore, begin by segmenting a surface using an automated graph-theoretic approach [21], which incorporates contextual information (as the segmentation is carried out in 3-D) and also ensures the global optimality of the segmented surface. The surface between the inner and outer segments of the photoreceptor cells (depicted in Fig. 4.3) is used, as it can be easily and reliably detected in OCT volumes. Two stages of smoothing thin-plate splines (described below) are then used to estimate the distinct artifacts seen in OCT images. In the first stage, a smoothing thin-plate spline is used to estimate and correct the tilt artifacts commonly seen in the B_f -scans. At this stage, the scleral curvature is also removed. A second spline-fit is then used to model and correct the rapidly varying motion artifact characteristic of the B_s -scans.

Thin-plate splines were first formulated by Duchon [71], who compared the process to the physical bending of a thin sheet of metal. The TPS formulation [72,73] used in our method generates a smoothing interpolation function as follows:

We begin by defining f to be a function that maps R^2 to R for a given set of N points in R^2 , $\nu = \{\nu_i: i = 1, 2, \dots, N\}$. Then, the thin plate spline interpolation function s will be one that minimizes the energy equation:

$$E_{tps}(s) = \sum_{i=1}^N \|s(\nu_i) - f(\nu_i)\| + \lambda \int_{R^2} \left(\frac{\partial^2 s(\theta)}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 s(\theta)}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 s(\theta)}{\partial y^2} \right)^2 d\theta, \quad (4.1)$$

where x and y are the two components of θ . Let us also define $[\phi_1(t), \phi_2(t), \phi_3(t)] = [1, x, y]$. Now, for the above energy equation there exists a unique minimizer s_λ given by

$$s_\lambda(\theta) = \sum_{k=1}^3 d_k \phi_k(\theta) + \sum_{i=1}^N c_i E(\theta - \nu_i), \quad E(r) = r^2 \log r \quad (4.2)$$

The parameters d_k and c_i can be estimated by solving N linear equations. Thus, a 2-D TPS can be fit to a 3-D segmented surface to obtain a smooth (where the smoothness is controlled by λ) 3-D reference plane with respect to which the dataset is flattened. Henceforth in this paper, we shall refer to the surface obtained through the 2-D TPS fit as the 3-D spline surface, and the curve obtained through the 1-D TPS fit as the 2-D spline curve.

To compensate for the two different artifacts seen in the dataset, the flattening is done in two steps (see Fig. 4.4):

1. A 2-D TPS is fit to the surface, where the number of control points used is determined by the size of the surface along each axial dimension. At this stage, the number is set to 10% and 5% of the dimensions along the x and y axial directions, respectively, and the control points are evenly distributed along each direction. A relatively large smoothing regularization term ($\lambda = 0.1$) is used so that the 3-D spline-fit surface thus created is relatively smooth and approximates the overall curvature of the retinal surfaces seen in the B_f -scans.

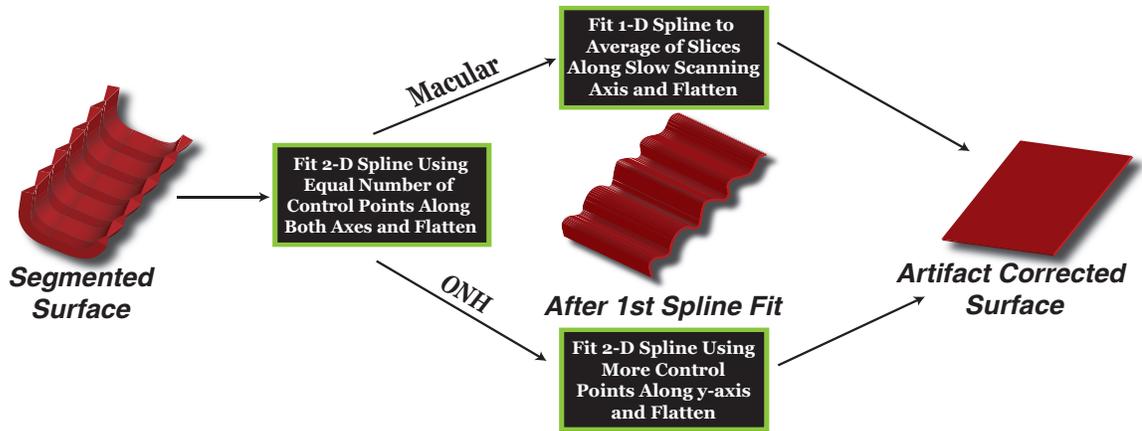


Figure 4.4: Overview of method to determine the flattening plane. The flattening plane is determined by fitting a spline to a surface twice, to eliminate the two distinct artifacts seen in these images.

(Experimentally, we found that values of λ between $[0.07, 0.13]$ provide consistent results.) However, at this stage, the large regularization term used makes it difficult to accurately estimate the artifacts in the slow-scanning direction (B_s -scans). The dataset is now flattened with respect to this reference plane by realigning the columns of the dataset, which eliminates the curvature seen in the B_f -scans as well as any tilt artifacts that may be present.

2. The rapid variations seen in the B_s -scans can be corrected in one of two ways:

- (a) The previously computed 3-D spline surface estimates and corrects the artifact along the fast-scanning axis, thus a single 2-D spline curve (computed using a 1-D TPS) can be used to model the artifacts in the B_s -scans. The segmented surface is averaged across the fast-scanning direction to create a single 2-D vector, which approximates the artifact seen in the B_s -scans. A single 1-D spline can now be fit to this vector using evenly spaced control points (totaling 25% of the axial dimension) and a relatively small regularization term ($\lambda = 0.07$) to give us a 2-D spline curve, which is used

to correct the motion artifact in all the B_s -scans. (Consistent results were obtained for $\lambda = [0.01, 0.07]$.) Note that the initial creation of the averaged 2-D vector (across all of the B_s -scans) helps to ensure that only the motion artifacts in this direction will be corrected rather than also flattening local retinal distortions, such as those from pathology. This method can be used to correct artifacts in macula scans as the volumes do not contain any regions where the surfaces are discontinuous.

- (b) In the case of optic nerve head (ONH) centered scans (where the surfaces become indistinct at the optic disc), a second 2-D thin-plate spline is fit to the new surface. The 2-D TPS now uses a larger number of control points in the y -direction (25% of the axial dimension) than the x -direction (5% of the axial dimension). These control points are chosen from outside the optic disc region, where the margin is approximated using a circle (2.1mm in diameter). Note that in the case of abnormal optic discs, this margin can be more precisely (and still automatically) determined [28]. The regularization term used in this step is also smaller ($\lambda = 0.05$), enabling the resulting 3-D spline surface to model the rapid variations seen along the slow scanning axis. (Consistent results were obtained for $\lambda = [0.03, 0.07]$.)

The final artifact-corrected dataset is now obtained by flattening the image once more with respect to the plane obtained in one of the two ways described above.

4.3 Experimental Methods

The availability of the paired-scans obtained with orthogonal scan patterns and the stereo fundus photographs, allowed for the method to be validated in two ways as described below.

4.3.1 Validation using paired OCT scans with orthogonal fast-scanning axes.

Since the artifacts seen in OCT images are strongly dependent on the orientation of the fast scanning axis, a pair of OCT scans acquired from the same eye with orthogonal fast scanning axes can be used to assess the accuracy of the artifact correction process. The tilt and low variation artifacts associated with B_f -scans now appear in perpendicular scans in the second dataset, and the same is true of high frequency variations associated with B_s -scans, as depicted previously in Fig. 4.2. The curvature and tilt artifacts associated with the B_f -scans are far easier to correct than the rapid variations seen in the B_s -scans, thus, the B_f -scans from one artifact corrected dataset can be used quantitatively to validate the ability of the proposed method to correct artifacts in the B_s -scans of the second dataset.

The quantitative measure of the accuracy of the artifact correction process is expressed using the mean unsigned difference in the location of a particular surface before and after the artifact correction process. The surface between the inner and outer segments of the photoreceptors was segmented and used in the spline-fit, and thus is available for use in the validation as well. An absolute comparison in microns is possible since this surface is flattened to the same depth in the z -axis.

The acquisition process of the OCT datasets creates volumes that are roughly rotated by 90° degrees with respect to each other. Thus, the datasets must first be registered to each other before any comparisons can be made. This was done by manually selecting two correspondence points in the 2-D projection images of the paired datasets. The projection images were created from a small number of slices near the segmented surface (between inner and outer segments of the photoreceptors), as the vessels are much clearer in projection images created in this manner [74]. Vessel crossings and bifurcations can be used as correspondence points. A rigid transformation (as expressed below) can now be used to align the B_f -scans of one dataset with the

B_s -scans of the second. It is easily apparent that two points are sufficient to compute the transformation matrix.

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta & \delta x \\ -\sin\theta & \cos\theta & \delta y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \quad (4.3)$$

In addition to the orthogonality of the scans, the images are also anisotropic in the x and y directions. In order to register the volumes better and make a true comparison between the segmented surface from the orthogonal scans, the projection images and the segmented surface are interpolated (along the smaller dimension) to make them isotropic. The mean unsigned difference between the segmented surface in both datasets can now be computed (within the common registered area) from the original, partially corrected (using a single 3-D spline surface) and the final artifact-corrected image.

4.3.2 Validation using 3-D reconstructions of the optic nerve head.

Tang *et al.* [75] reported a method for the reconstruction of the shape of the ONH from stereo fundus photographs. The 3-D shape estimate is obtained by finding corresponding pixels from two stereo images of the ONH, taken from two slightly different angles. The two image planes are known to be horizontally displaced, but can be assumed to be co-planar. Since the horizontal disparity is known to be inversely proportional to the depth associated with the 3-D pixel, a depth map can be created using pixel correspondences. The depth maps thus created (Fig. 4.5(b)) show the shape of the retina at the ONH region, and since they are created from fundus photographs they are free of the axial artifacts associated with OCT scans.

The structure obtained from the OCT images that is most comparable to the depth maps is the location of the inner limiting membrane (ILM). Before any comparison can

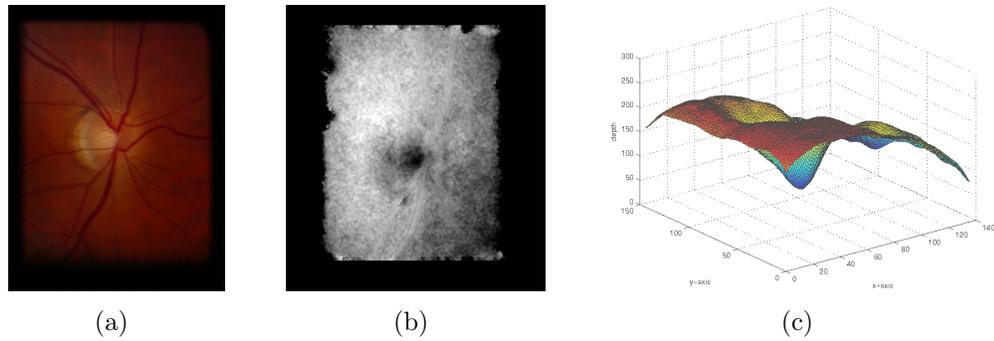


Figure 4.5: Fundus photograph and its corresponding disparity maps. (a) One of a pair of stereo fundus photographs of a glaucomatous eye. (b) The disparity map constructed from the stereo fundus photographs. (c) The smoothed disparity map in 3-D showing the overall shape of the optic nerve head.

be made between the depth maps and the location of the ILM, we have to compensate for three important characteristics of stereo fundus photographs.

1. The fundus photographs are not in the same reference frame as the OCT images, thus the depth maps must first be registered to the OCT dataset. Vessel locations were used to guide the rigid registration of the fundus images to the 2-D projection image created from the OCT image, as described in Section 4.3.1.
2. The depth from stereo estimations contain noise, which is seen in the depth maps. Thus, the depth maps must first be smoothed to validate the artifact correction process. The smoothing was done so that the noise was suppressed while the shape information from the depth maps was retained. Fig. 4.5 shows the fundus photograph, its corresponding depth map and 3-D rendering of a smoothed depth map.
3. The depth maps do not provide quantitative depth information, as these are serial stereo photos and stereotactic positions (angle between camera positions) is not available. Thus, the location of the segmented surface from the OCT

images must be scaled and expressed in normalized units. For this, we consider the depth of the surface from the top of the OCT dataset and normalize this depth by dividing by 200. We then scale this normalized depth to match the scale of the depth maps. The z -axis depth location of the reference plane for all datasets is maintained at the same value to minimize variations between the datasets.

A pixel by pixel comparison can now be made between the smoothed depth maps and the normalized depth of the ILM in the flattened OCT datasets.

4.4 Modeling the Shape of the Eye

The pairs of OCT datasets acquired from the same eye can be used to estimate the shape of the eye since the B_f -scans, which contain information about the retinal curvature, are available along two perpendicular axis. Thus, in datasets with small or no tilts induced by the incorrect positioning of the camera, the B_f -scans from one dataset can be used to “correct” the B_s -scans in the second dataset, creating a dataset where the retinal surfaces now reflect a better approximation of the shape of the eye.

A 2-D TPS is fit to the segmented surface from each of the datasets to create a 3-D interpolated isotropic surface. The 3-D spline surface thus created is not only smooth and isotropic, but is also less dependent on the segmentation result. Figs. 4.6(a) and (b) show the isotropic surfaces created from the datasets with the horizontal and vertical fast scanning axes, respectively. The B_f -scans from the dataset with the vertical fast scanning axis can now be used to correct the artifacts in the B_s -scans of the dataset with the horizontal fast scanning axis. Since the datasets are acquired from the same area of the retina, the isotropic surfaces can now be used to estimate the z -axis translations required to correct the rapid variations seen along the slow scanning axis. Correcting the artifact in this manner will retain the ocular shape of

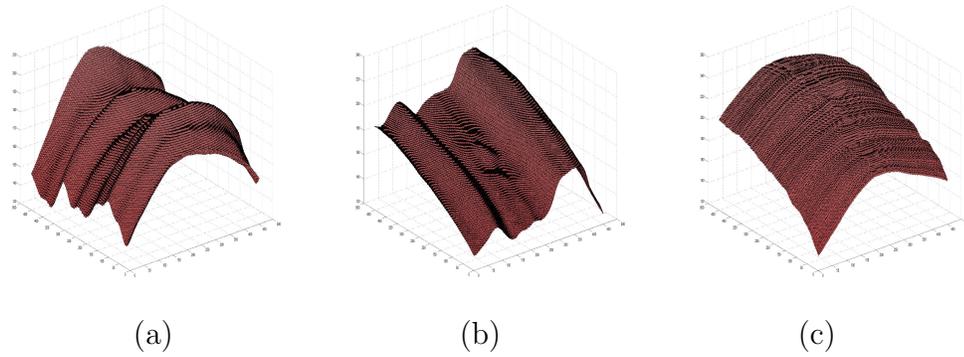


Figure 4.6: Estimating the shape of the eye from paired OCT datasets with orthogonal fast scanning axes. The isotropic surfaces obtained by fitting a thin-plate spline to a segmented surface from (a) the dataset with the horizontal fast scanning axes, and (b) the dataset with the vertical fast scanning axes. (c) The estimate of the “true” curvature of the retina.

the retina. The final shape corrected surface is as shown in Fig. 4.6(c).

4.5 Results

The quantitative results are summarized in Table 1. In the first validation technique (using the paired macula-centered scans), a segmented surface from the two datasets was compared before and after the artifact-correction process. The mean unsigned differences seen in the original and partially corrected (single spline-fit) datasets were $98.97 \pm 39.45 \mu\text{m}$ and $23.36 \pm 4.04 \mu\text{m}$, respectively. The artifact-corrected datasets, on the other hand, only showed a mean unsigned difference of $5.94 \pm 1.10 \mu\text{m}$, which is significantly smaller ($p < 0.001$, from two-tailed paired t -test) when compared to the partially corrected datasets. Fig. 4.7 shows an example of the 3-D representation of the original surface used to create the reference plane, the surface after the first spline-fit and the surface in the final artifact-corrected image.

Fig. 4.7(b) shows the segmented surface in a macula-centered OCT dataset after the first spline fit, and it is easy to see the periodic nature of the artifact along the y -axis. The single 2-D spline curve computed (using a 1-D TPS fit to an averaged 2-D

Table 4.1: Mean unsigned difference seen in the original, partially corrected and final artifact-corrected images.

	Original Volume	Single Spline Fit	Artifact-Corrected Volume
Validation 1 [◊]	98.97 ± 39.45	23.36 ± 4.04	5.94 ± 1.10
Validation 2 [*]	0.302 ± 0.134	0.142 ± 0.036	0.134 ± 0.035

[◊] Mean unsigned difference was computed between the location of a segmented surface in the paired OCT datasets and expressed in μm .

^{*} Mean unsigned difference was computed between the disparity maps and the normalized depth of the ILM. Average value is expressed as mean \pm standard deviation in normalized units. Differences were computed in regions where the disparity maps were well defined. The optic disc was avoided.

vector) along this axis is more than sufficient to estimate this artifact and eliminate it, as can be seen in Fig. 4.7(c). Figs. 4.8(a) and (b) shows the central B_f -scan and B_s -scan, respectively, from a macula-centered OCT dataset before the artifact correction process. The same slices after the artifact correction process are depicted in Figs. 4.8(c) and (d).

In the second validation method, 3-D depth maps created from stereo fundus images were compared to the ILM. As the depth maps are created from the fundus photographs and can sometimes show only a small region around the optic disc, care was taken to ensure that the difference was only calculated in areas where the disparity map was well defined. The optic disc region was also avoided. The mean unsigned difference (computed in normalized units) seen in the original and partially corrected (single spline fit) datasets was found to be 0.321 ± 0.134 and 0.142 ± 0.036 , respectively. The artifact-corrected datasets showed a mean unsigned difference of 0.134 ± 0.035 , which is significantly smaller ($p < 0.001$, from two-tailed paired t -test) when compared to the original datasets.

Figs. 4.7(d), (e) and (f) show the segmented surface from an ONH centered dataset in the original dataset, the dataset after correction with a single 3-D spline surface, and the dataset after the second 3-D spline surface correction process, respectively.

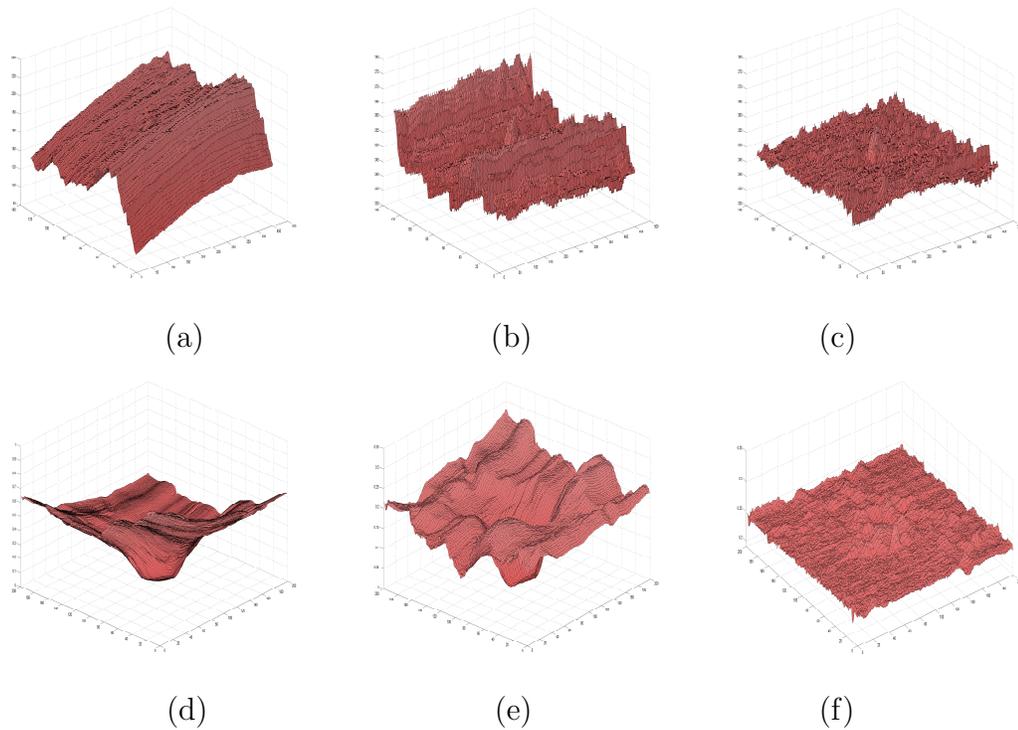


Figure 4.7: Examples of the surface used to create the reference plane. (a), (b) and (c) show the segmented surfaces from a macula-centered OCT dataset in the original, partially corrected and final artifact-corrected image, respectively. (d), (e) and (f) show the segmented surface from an ONH-centered OCT dataset in the original, partially corrected and final artifact-corrected image, respectively.

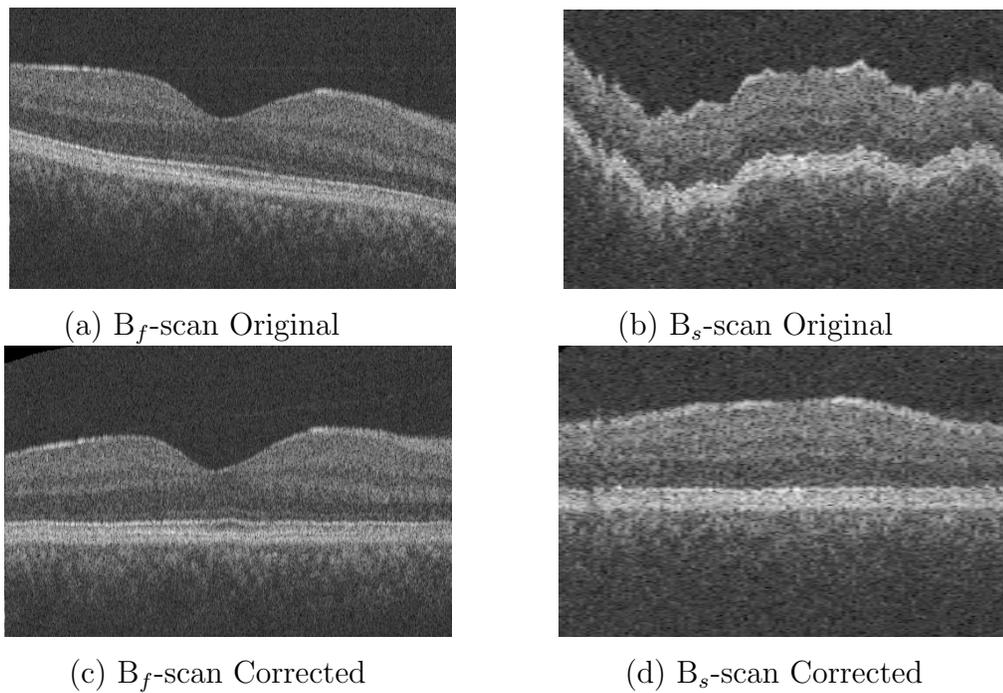


Figure 4.8: Central B_f -Scan and B_s -scan slices from an OCT dataset before and after flattening. Original (a) B_f -Scan and, (b) B_s -Scan before artifact correction, respectively. Artifact corrected (c) B_f -Scan and (d) B_s -Scan, respectively.

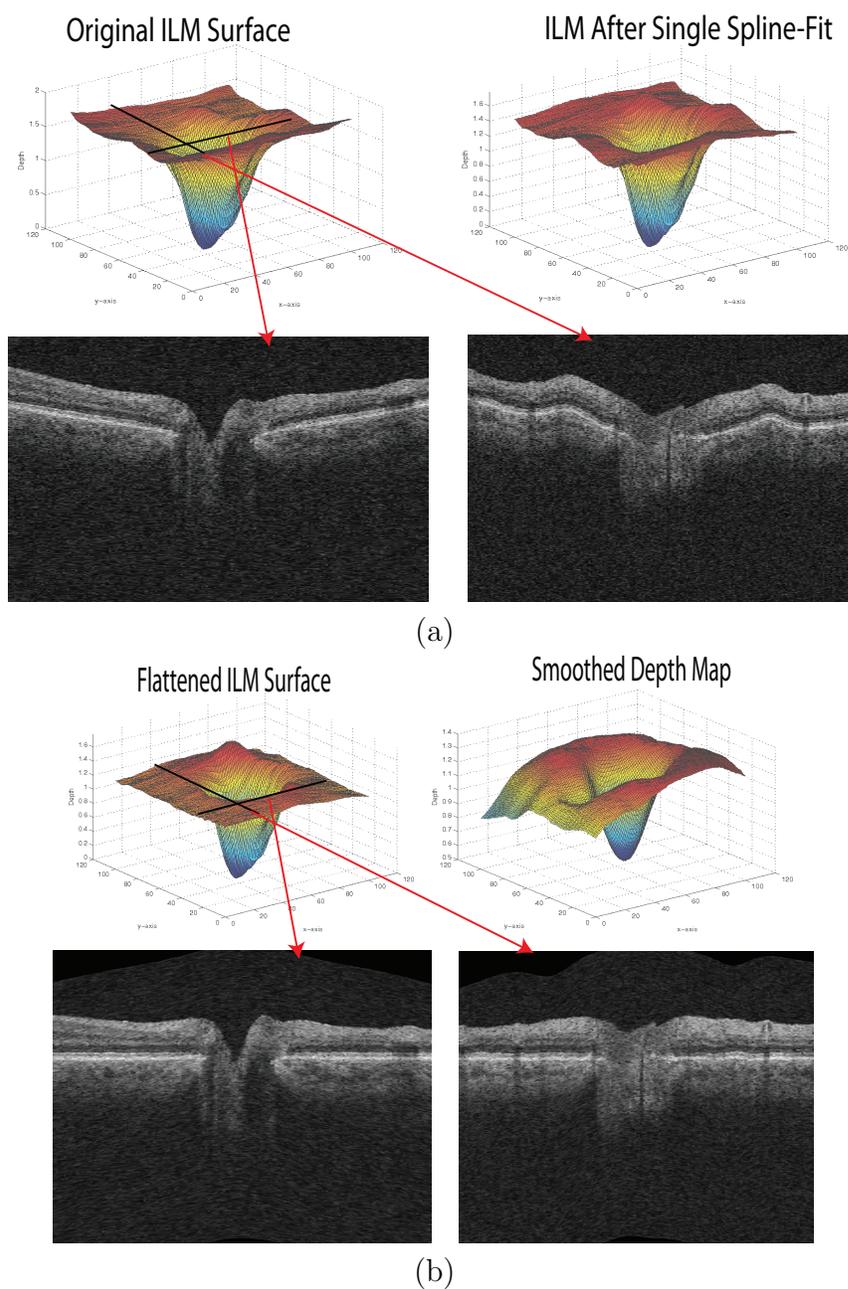


Figure 4.9: An ONH-centered dataset before and after the axial artifact correction. (a) The ILM in the original dataset and after the first spline fit. B_f -scan and B_s -scan from the original dataset from the locations as indicated by arrows in red. (b) The ILM in the flattened artifact corrected dataset with the same B_f -scan and B_s -scans from the flattened dataset. Smoothed depth image also included.

The optic disc disrupts the surfaces (as seen in Figs. 4.7(d) and (e)) necessitating the use of a second 3-D spline surface instead of a 2-D spline curve. Fig. 4.9 shows the ILM and two slices from an ONH-centered dataset in different stages of artifact correction. The vast difference between the original and artifact-corrected images is easily apparent in the slices depicted in Figs. 4.9(a) and (b), respectively.

4.6 Discussion & Conclusion

Eliminating motion artifacts in OCT images is important as it brings the dataset into a consistent shape and makes visualization and subsequent analysis easier. While the removal of the scleral curvature seen in the B_f -scans is undisable in some applications, numerous applications such as the automated segmentation of intraretinal layers [17, 21], and the volumetric registration of OCT-to-OCT [59] datasets would benefit from the consistency of an artifact-corrected dataset.

In this work, we have presented a method for the correction of axial artifacts using thin-plates splines that estimate the distortions along the fast and slow scanning axes in OCT retinal scans. Our results show that the TPS approach is effective, as it is able to create a globally smooth surface whose properties are easily controlled by the regularization parameter and the number of control points used. The proposed method aims to eliminate all of the artifacts, and therefore does not retain any information about the shape of the retina; however, given the new availability of orthogonal OCT scans on clinical OCT machines, the shape can be estimated and reconstructed. It is also important to note that the approach does not alter the data in any way other than the z -axis translation (which is a reversible transformation), and thus does not affect any measurements derived from the A-scans. The time complexity of the method largely depends on the number of control points used in the TPS fit, as the surface segmentation takes under a minute. Our implementation (in C++, on a 2.8GHz six-core AMD Opteron processor) took a total time of 7-9 minutes to segment the surface and correct the artifacts in the OCT volumes.

The robustness of the correction procedure is evident from the results obtained on a diseased set of OCT scans, where the surface segmentation is more prone to error. While the dependence on the segmentation result is undesirable, it does provide vital information in scans where the camera has been incorrectly positioned and the retinal surfaces appear skewed. In such situations, the 3-D spline surface flattening procedure would show better results than 2-D rigid registrations methods, as the 2-D registration approaches only aim to correct the motion artifacts, but do not address the tilt artifacts. Furthermore, the 2-D registration could introduce artifacts into the image in the form of unwanted translational or rotational changes, as they are not guided by 3-D contextual information. The use of SLO images [65] or fundus photography would be necessary to ensure that such artifacts do not affect the end result. Alternatively, a combination of segmentation and registration-based methods could be used to retain 3-D contextual information during the artifact correction process.

Thus, in summary, the proposed two-stage flattening approach, is able to correct multiple types of axial artifacts (some for which the initial 3-D surface segmentation is necessary or useful for artifact-modeling purposes), while still demonstrating a robustness against any small local disruptions or errors in the initial segmentation result. The availability of orthogonal scans allows for the correction of artifacts as well as the reconstruction of the true shape of the retina.

CHAPTER 5

INCORPORATION OF TEXTURE-BASED FEATURES IN AN OPTIMAL GRAPH-THEORETIC APPROACH WITH APPLICATION TO THE 3-D SEGMENTATION OF INTRARETINAL SURFACES IN SD-OCT VOLUMES

As described in Section 3, the graph-theoretic approach simultaneously segments multiple surfaces in n -D while ensuring the global optimality of the solution with respect to a cost function. Thus, the careful design of the cost function is an important step. This process however, is far from trivial and can be time-consuming, particularly when the image quality is affected by disease. Fig. 5.1(a) shows a B_f -scan from an image obtained from a subject that presented with glaucoma. Not only are the retinal layers difficult to discern, but some have also been thinned by the disease. A baseline approach [27] that did not incorporate any learned information performed adequately when utilized for the segmentation of glaucomatous scans and showed errors that were comparable with the inter-observer variability. However, noticeable errors were often seen at the optic nerve head (see Fig. 5.1(c)) and near large vessel shadows (see Fig. 5.1(f)), especially when a number of vessels are close to each other. In such regions the inclusion of learned textural features in the cost function [29] in lieu of relying on gradients (that are quite small) could help prevent such errors.

When segmenting multiple regions, similarities could exist between two or more regions as is the case in retinal OCT images. As shown in Fig. 5.2(b), the retinal nerve fiber layer (RNFL) is distinctively bright, while the inner plexiform and outer plexiform layers are of medium intensity and the inner and outer nuclear layers have low intensity. In standard classification-based approaches [11, 76], the goal is to find the best classification possible for each region of interest and requires each region to be accurately distinguished from all other regions. This may however, be unnecessary when combining such an approach with the graph-theoretic approach as this method only requires that regions be distinct from their immediate neighbors rather than every other region in the image. This could allow for the regions to be grouped on

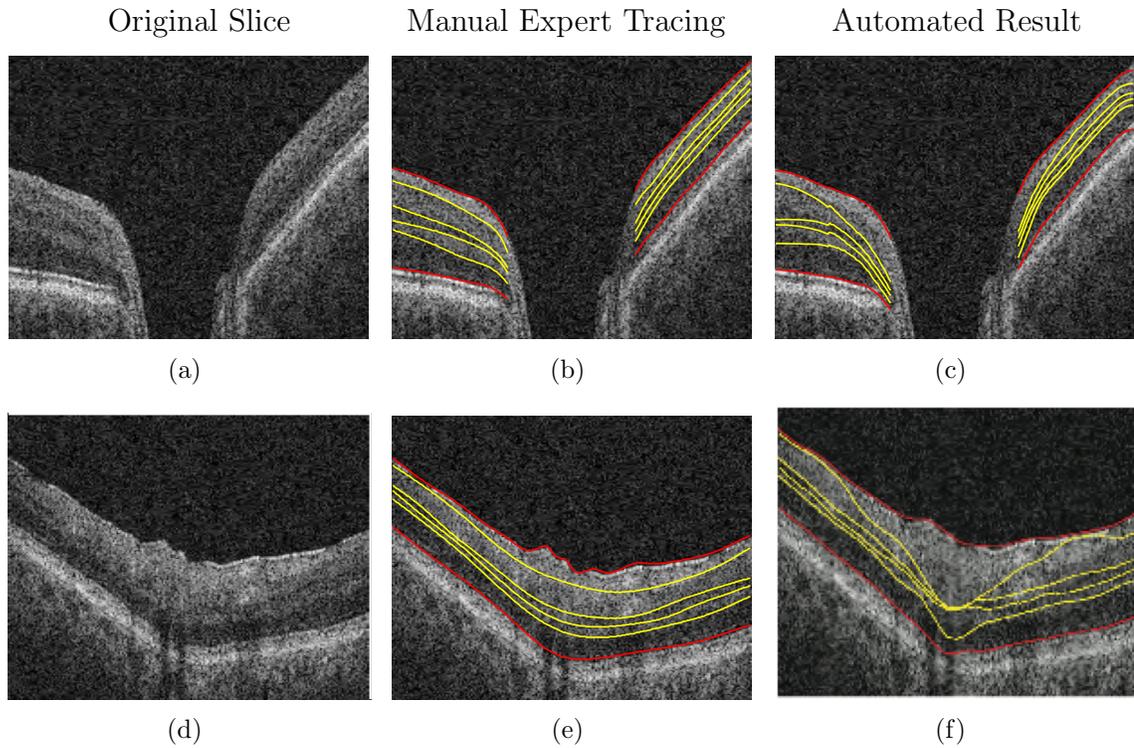


Figure 5.1: Two example of diseased scans that proved problematic for the baseline approach that did not incorporate any learned features. (a) A central xz -slice from an optic nerve head centered glaucomatous scan zoomed in on the inner retinal surfaces. (b) The same slice with the manual tracings overlaid. (c) The results obtained using the baseline approach. Another example (d) slice from a glaucomatous scan, (e) with the manual tracing overlaid and the (f) automated result.

the basis of similar properties.

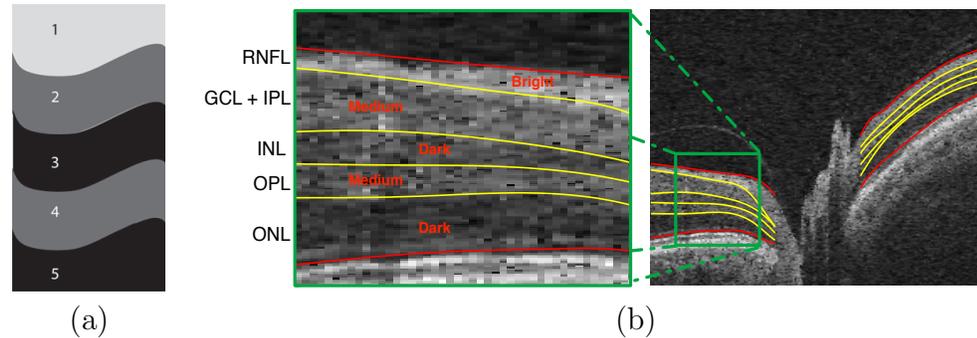


Figure 5.2: Illustration of the interacting surfaces seen in SD-OCT retinal images. (a) A generalized depiction of the multi-layer problem, where one or more layers have similar intensity. (b) An example of an OCT volume from which the retinal surfaces are segmented.

A similar approach can be used to design the on-surface cost term as well as to incorporate learned shape priors into the graph-theoretic method as described in Chapters 7 and 8, respectively. In this chapter, however, we focus on the incorporation of learned textural features into the in-region cost term of the cost function. The method is utilized for the segmentation of six retinal surfaces in 10 SD-OCT scans obtained from patients diagnosed with glaucoma. A variety of textural features in 2-D as well as 3-D were learned from a training set and incorporated into the cost function utilized by the graph-theoretic method. Additionally, the impact of grouping regions on the basis of a similar property, such as intensity, is also investigated.

5.1 Methods

In order to design the cost functions, we must first find descriptive features for the regions of interest. For this, a number of textural descriptors were investigated including intensity-based features (mean intensity, variance, entropy, kurtosis and skewness), co-occurrence matrix features (energy, entropy, contrast, correlation and

inverse difference moment) [77], run-length features (short run, long run, gray level non-uniformity, run length non-uniformity and run percentage) [78] in addition to features generated using a variety of filter banks such as Haar [79], Gabor [80], and steerable-Gaussians.

For the Haar features, kernel sizes $w = \{1, 2, 4, 6\}$ were used to create 2-rectangle and 4-rectangle features. Gabor filters, which have found widespread use in texture analysis and classification, are not only linear and local, their orientation and spatial frequencies can be tuned to meet specific requirements. Essentially, they act as band pass filters with optimal localization properties in the spatial and frequency domain. The 2-D Gabor filter is represented as:

$$g(x, y)_{(\sigma, \gamma, \lambda, \theta, \psi)} = e^{-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}} e^{i(2\pi \frac{x'}{\lambda} + \psi)}, \quad (5.1)$$

$$\text{where } x' = x \cos \theta + y \sin \theta, \quad y' = -x \sin \theta + y \cos \theta. \quad (5.2)$$

The parameters $\sigma, \gamma, \lambda, \theta$ and ψ are the standard deviation, spatial aspect ratio, wavelength, orientation and the phase offset, respectively. A Gabor filter bank [81] was created using the symmetric ($\psi = 0^\circ$) and anti-symmetric ($\psi = 90^\circ$) filter responses of the Gabor filter with standard deviation $\sigma = \{1, 2, 4, 6\}$, and angle $\theta = \{-90^\circ, -60^\circ, -30^\circ, 0^\circ, 30^\circ, 60^\circ, 90^\circ\}$. The wavelength (λ) was set to $\sigma/0.56$ as this corresponds to a half-response spatial frequency bandwidth of one octave and the spatial aspect ratio (γ) was set to 1. The energy features [80] were also computed by summing the squared symmetric and anti-symmetric filter responses for particular values of σ and θ , as shown below:

$$E_{\sigma, \theta}(x, y) = \sqrt{G(x, y)_{(\sigma, \theta, 0^\circ)}^2 + G(x, y)_{(\sigma, \theta, 90^\circ)}^2} \quad \text{where,} \quad (5.3)$$

$$G(x, y)_{(\sigma, \theta, \psi)} = \sum_{\eta=0}^N \sum_{\nu=0}^M I(\eta, \nu) g_{(\sigma, \theta, \psi)}(x - \eta, y - \nu). \quad (5.4)$$

The Gaussian filter bank [82] consisted of oriented filter derivatives up to order 1 at scales $\sigma = \{1, 2, 3, 4, 5\}$ and orientations $\theta = \{0^\circ, 30^\circ, 60^\circ, 90^\circ, -60^\circ, -30^\circ, -90^\circ\}$, where

$$G_1^\theta = \cos(\theta)G_1^0 + \sin(\theta)G_1^{90}, \quad \text{where} \quad (5.5)$$

$$G_1^0 = \frac{\partial}{\partial x}(e^{-(x^2+z^2)}) \quad \text{and} \quad G_1^{90} = \frac{\partial}{\partial z}(e^{-(x^2+z^2)}). \quad (5.6)$$

These features were computed in 2-D within each B-scan using Eq. 9.11 and also in 3-D where

$$G_1^0 = \frac{\partial}{\partial x}(e^{-(x^2+y^2+z^2)}) \quad \text{and} \quad G_1^{90} = \frac{\partial}{\partial z}(e^{-(x^2+y^2+z^2)}). \quad (5.7)$$

A classifier trained on these features can now be used to create probability maps for the regions of interest. k -Nearest neighbor (k -NN) classifiers are relatively simple non-parametric classifiers that assign the input feature the same label as its k nearest neighbors. Regression analysis can be performed by computing the ratio n_r/k , where n_r is the number of neighbors belonging to a particular class r and k is the total number of neighbors used in the k -NN classifier. This classifier is simple, yet effective. Its performance however, is heavily influenced by the dimensionality of the feature space. Thus, to reduce computational time, the sequential forward floating selection (SFFS) algorithm [83] was used to select the most descriptive features that maximized the accuracy and specificity of the classification result. This feature selection process can be a time-consuming task and therefore, a GPU-based implementation was used that significantly reduced the computational time.

The overview of the segmentation approach is depicted in Fig. 5.3. This method is quite similar to the previously described approach [27] that was utilized to segment intra-retinal layers in ONH-centered SD-OCT images, with the only addition being the learned cost function. In the case of the 2-D features the probability maps

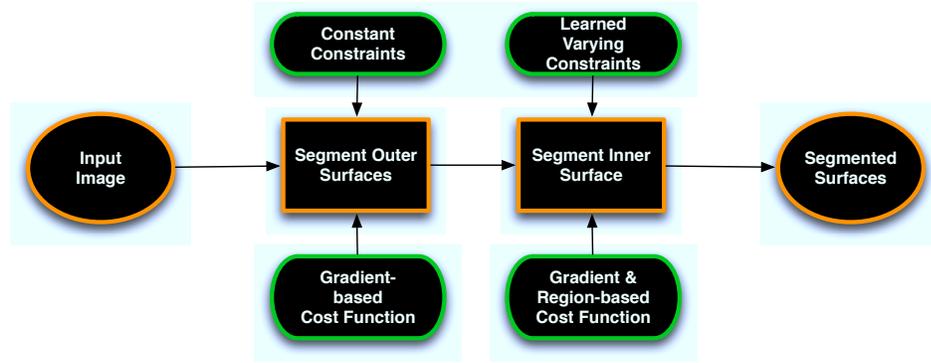


Figure 5.3: A schematic sketch of the proposed segmentation method that incorporated learned textural features into the cost function utilized by the graph-theoretic method.

are generated in a slice by slice fashion to create a 3-D volumetric cost function. The outer surfaces (the inner limiting membrane (ILM) and the IS/OS line and the Bruch's membrane) are segmented first, followed by the segmentation of the four inner surfaces. As the layers in the outer retina are quite thin, learning region-based textural information here is difficult. Thus, the impact of the learned textural features is investigated solely with respect to the inner retinal layers. As the graph-optimization aims to find the minimum-cost solution, the probability maps were inverted so that the voxel values reflect the unlikelihood of that voxel belonging to a particular region [21].

5.2 Experimental Methods

The data used to train and test the texture-based regional cost functions consisted of 10 optic nerve head (ONH) OCT volumes from 10 patients diagnosed with glaucoma, acquired on spectral-domain Cirrus (Carl Zeiss Meditec, Inc., Dublin CA). Each dataset had dimensions of $6\text{mm} \times 6\text{mm} \times 2\text{mm}$ and $200 \times 200 \times 1024$ voxels. Manual tracings from two independent observers were obtained from 10 slices selected randomly (from 10 sections of 20 slices each) in each dataset. The texture-based re-

gional cost functions were generated in three different ways:

1. 2-D texture features where
 - the regions were not grouped (Approach I), and
 - the regions were grouped on the basis of intensity (Approach II).
2. 3-D texture features where the regions were grouped on the basis of intensity (Approach III).

A cross-validation scheme was used where the classifier was trained on eight datasets (80 slices), while the other two datasets were used as the test set. Five separate experiments were conducted in this manner and the average unsigned border position error was computed between the automated result and the average of the manual tracings.

5.3 Results

The probability maps obtained for Approaches I, II and III are as shown in Fig. 5.4. The manual tracings from one independent observer are also shown on the images to indicate the boundaries of the regions. The grouping of the regions on the basis of intensity reduced the misclassifications seen in Figs. 5.4(c) and (e) as shown in Fig. 5.4(h).

The average unsigned border position errors noted per surface is as tabulated in Table 5.1. The inter-observer variability, while not in the table [27] for surfaces 2, 3, 4 and 5 were 9.92 ± 2.65 , 7.48 ± 1.89 , 7.07 ± 1.31 and 7.04 ± 1.31 , respectively. The overall inter-observer variability for the six surfaces was $8.85 \pm 3.85 \mu\text{m}$. Approach II, that utilized 2-D texture features and grouped the regions on the basis of intensity showed the smallest unsigned border position error. A *t*-test was used to statistically compare the results obtained using the baseline approach and Approaches I, II and III. While the errors noted for surfaces 4 and 5 reduced when using Approach I, the

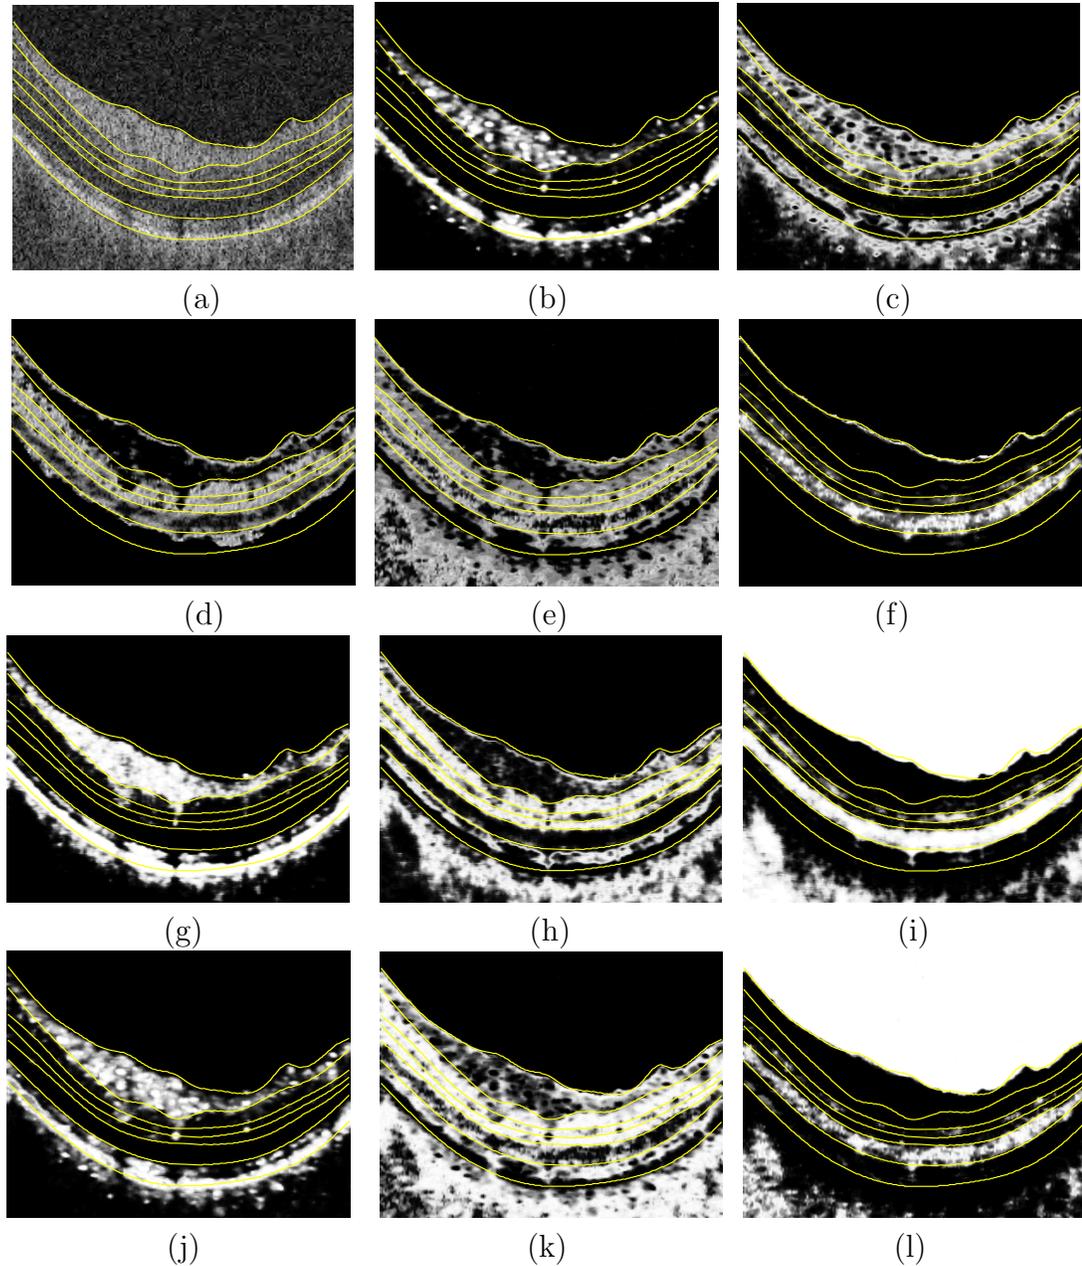


Figure 5.4: Probability maps obtained for the various region with and without grouping on the basis of similar properties. (a) A slice from an OCT dataset showing the manual tracings from one independent observer. 2-D features were classified to create soft labeled probability maps for the (b) RNFL, (c) GCL + IPL, (d) INL, (e) OPL and (f) ONL layers, respectively, where the regions are assumed to be independent, and were therefore, not grouped. The probability maps obtained for the (g) high-, (h) medium- and (i) low-intensity categories using the 2-D features and the probability maps obtained for the (j) high-, (k) medium- and (l) low-intensity categories using the 3-D features. The tracings from one independent observer are overlaid on the images to indicate the true boundaries of the regions.

Table 5.1: Average mean unsigned border positioning error[†] computed on 10 glaucomatous datasets.

Surface	Baseline vs. Avg. Obs.	Approach I vs. Avg. Obs.	Approach II vs. Avg. Obs.	Approach III vs. Avg. Obs.
1	4.70 ± 1.59	4.70 ± 1.59	4.70 ± 1.59	4.70 ± 1.59
2	14.43 ± 5.63	16.71 ± 5.94	11.03 ± 3.87	11.07 ± 3.79
3	10.96 ± 4.06	15.87 ± 4.74	6.95 ± 2.08	7.42 ± 2.21
4	10.46 ± 2.79	8.63 ± 2.72	6.05 ± 1.95	5.70 ± 1.78
5	10.73 ± 2.78	7.09 ± 2.34	6.26 ± 2.04	6.60 ± 2.15
6	3.87 ± 1.32	3.87 ± 1.32	3.87 ± 1.32	3.87 ± 1.32
Mean	9.25 ± 4.03	9.48 ± 5.55	6.48 ± 2.50	6.56 ± 2.55

[†] Mean ± SD in μm .

overall results were not significantly different from those obtained using the baseline approach. Moreover, the errors noted in surfaces 2 and 3 when using Approach I were significantly ($p < 0.01$) worse than those obtained using the baseline approach. Approaches II and III, showed significant ($p < 0.001$) improvements over the baseline approach for each of the four surfaces (2 - 5) as well as in the overall segmentation result. The results obtained using these two approaches however, did not significantly differ from each other.

Fig. 5.5 shows the comparison of the segmentations obtained with and without the inclusion of texture, and it is easy to see that the 2-D textural features help the segmentation method the most. Fig. 5.5(a),(b) and (c) show the manual tracings from one independent observer overlaid on the OCT slices. The results without the inclusion of texture, after the inclusion of the 2-D features using Approach I and II and the 3-D features (Approach III) are as depicted in the second, third, fourth and fifth rows of images, respectively.

Fig. 5.6 shows the comparison of the segmentations obtained with and without the inclusion of texture on a second dataset. Here, we can see that the inclusion of texture in general improves the segmentation, especially in areas close to the optic

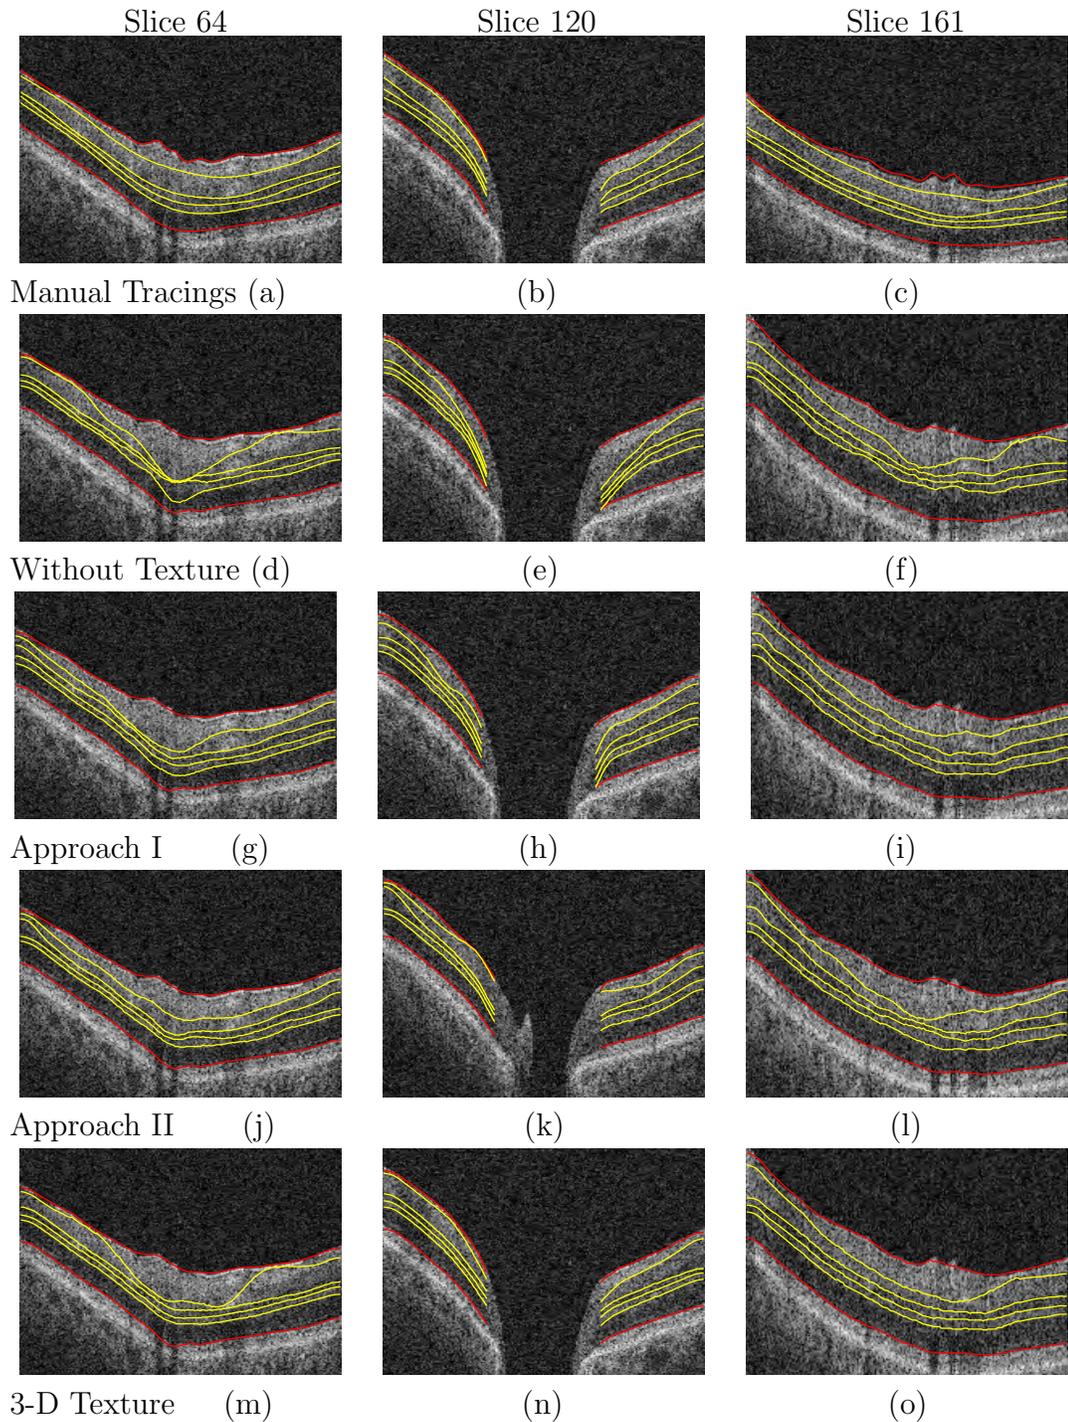


Figure 5.5: Segmentation results. Example slices (a)-(c) from an OCT volume showing the manual tracings from one independent observer, the segmentation result obtained (d) - (f) without the incorporation of texture, the incorporation of 2-D features obtained using (g) - (i) Approach I, (j) - (l) Approach II and (m) - (o) 3-D texture features, respectively.

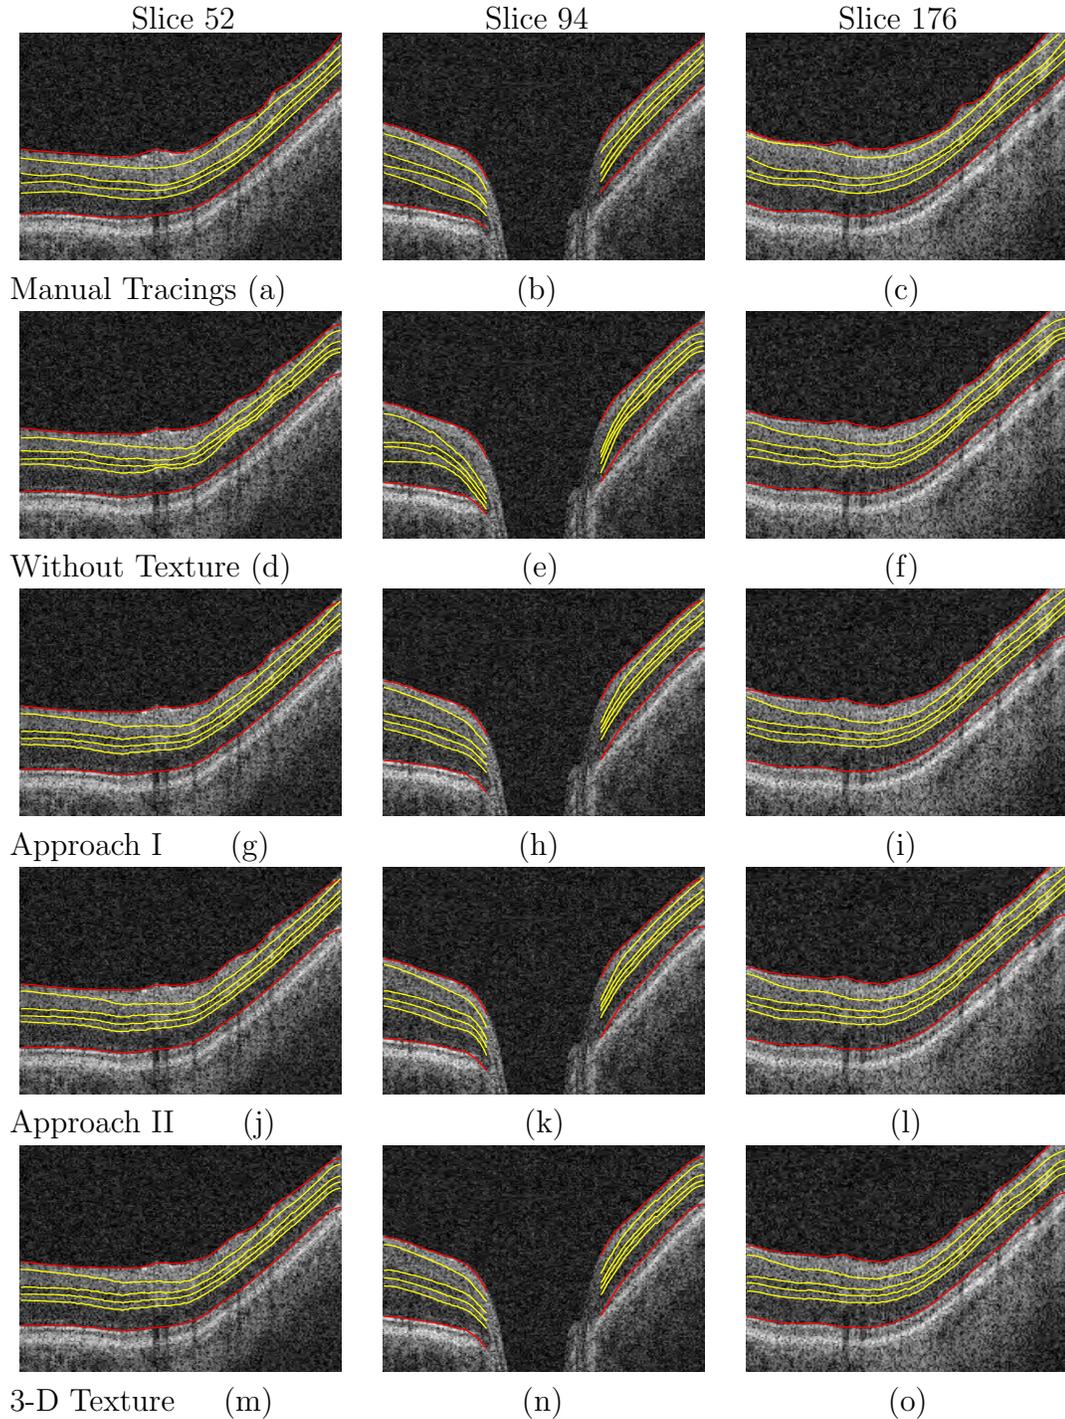


Figure 5.6: Segmentation results. Example slices (a)-(c) from an OCT volume showing the manual tracings from one independent observer, the segmentation result obtained (d) - (f) without the incorporation of texture, the incorporation of 2-D features obtained using (g) - (i) Approach I, (j) - (l) Approach II and (m) - (o) 3-D texture features, respectively.

nerve head. In Fig. 5.6 (e), which corresponds to the results obtained without incorporation textural features, the errors are large, but when texture is incorporated, the errors are significantly reduced.

5.4 Discussion

From the probability maps generated, we see that the segmentation method only requires that each region be separable from its immediate neighbors, thus the results from Approaches II & III were significantly better than those obtained using Approach I. Furthermore, the influence of the cost functions is apparent from the results obtained from Approaches I & II, and it is clear that a poorly designed cost function could adversely affect the overall segmentation accuracy. When the regions are not grouped, the classification is prone to misclassification errors between similar regions, such as the dark inner nuclear layer (INL) and outer nuclear layer (ONL) (see Figs. 5.4(d) and (f)).

The results also showed that Approach II (which uses 2-D features along with the grouped-region classification) produced the best results, fairing a little better than the 3-D features. The lack of improvement over the 2-D features and the significantly large computational time required for the 3-D feature generation has discouraged their further use.

In conclusion, we see that the incorporation of texture can improve the overall segmentation accuracy significantly. It must also be noted that the scans used come from diseased eyes, resulting in a higher degree of variability than would be seen in normal subjects, which is reflected in the high inter-observer variability. But despite the low quality of the images, we saw a significant improvement in accuracy in comparison to our previously reported results.

CHAPTER 6

AUTOMATED 3-D SEGMENTATION OF RETINAL SURFACES IN SD-OCT VOLUMES NORMAL AND DIABETIC MICE

The ability to discern and quantitatively measure retinal structures of interest in SD-OCT images of animal retinæ [84–86], such as mice [87,88] has led to the increasing use of this imaging modality in research studies in addition to its clinical applications. The automated segmentation of retinal layers in human scans [2–13,17,19] is a well-studied problem, however the same is not true of animal scans. Mishra et al. [14] and Yazdanpanah et al. [15] have previously described automated methods for the segmentation of retinal layers in SD-OCT images of rat retinas. However, the evaluation of these methods was limited to normal rats and the methods do not incorporate the three-dimensional contextual information available in volumetric scans. This is of particular importance given the increasing ability to obtain isotropic images from animal retinas that are rich in contextual information. The reproducibility of measurements from mice scans has been described briefly [89], however this experiment was limited to the reproducibility of total retinal thickness.

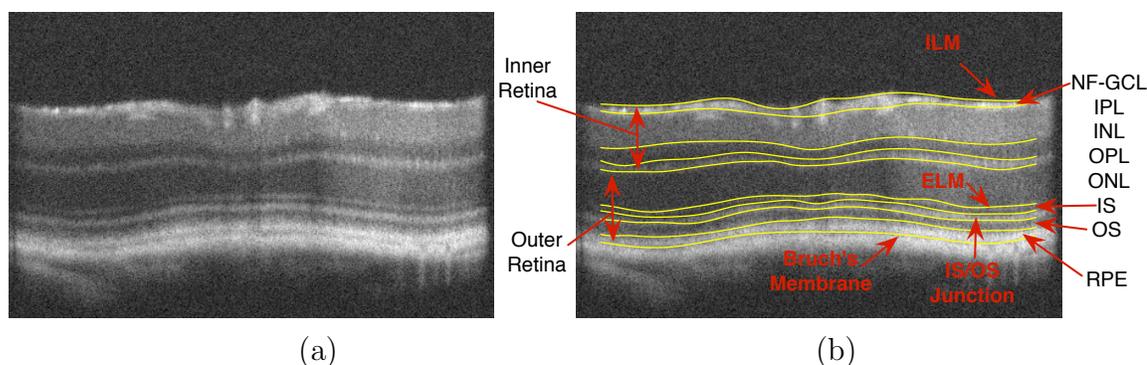


Figure 6.1: Slice from an SD-OCT volume showing 10 retinal surfaces. (a) A B-scan from an SD-OCT volume of a mouse eye with (b) the manual tracings of 10 retinal surfaces overlaid.

In this chapter, we detail a graph-theoretic method for the segmentation of the 10 retinal surfaces (see Fig. 6.1) in mice SD-OCT volumes. This method is an adaptation of a graph-theoretic approach that was previously [21, 25, 27] used of the automated segmentation of intra-retinal layers in human SD-OCT volumes. The accuracy of the method was assessed by comparing the segmentations to manual tracings provided by a retinal specialist. The reproducibility of thickness measurements was also assessed in scans obtained from normal and diabetic mice. This method did not incorporate any learned features, such as those presented in the previous chapter, and served as a baseline method which was used in subsequent experiments (described in Chapters 7 and 8).

6.1 Method

The graph-search approach for the segmentation of images, as described in our enabling work [25, 27], is a multi-resolution approach that begins with the segmentation of the outer surfaces (see Fig. 6.2), namely, ILM, ELM, the junction between the inner and outer segments (IS/OS line) and the bounding surfaces of the RPE. The inner retinal layers were segmented next, namely the nerve fiber-ganglion cell (NF+GC) complex, the IPL, the INL and the OPL. While the overall method remained the same, some changes were made in order to accommodate the differences in anatomy between human and mice retinas.

The outer surfaces segmentation used a 4-step multi-resolution approach as illustrated in Fig. 6.3, consisting of volumes downsampled by factors of 8, 4, 2 and the original volume itself. First, the ILM, the IS/OS junction and the BM were simultaneously segmented in a volume down sampled by a factor of 8. The results of the surfaces at this stage were used to create regions of interest in the higher resolution volumes. The ILM was segmented as a single surface in the higher resolution volumes, while the other outer surfaces were segmented simultaneously. Specifically, the IS/OS junction, the bottom of the OS and the bounding surfaces of the RPE were segmented

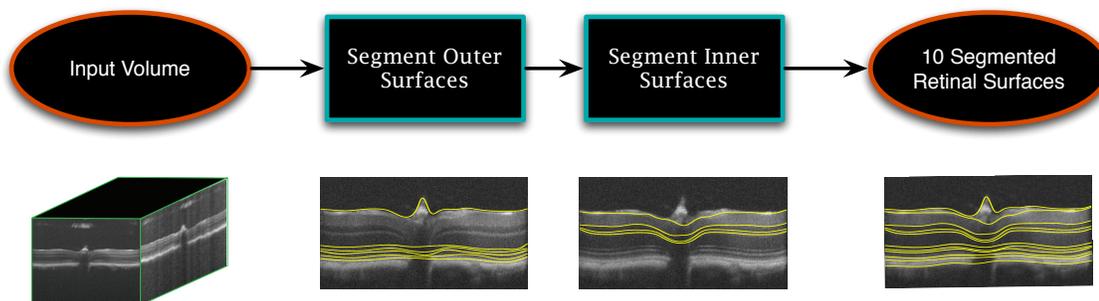


Figure 6.2: Overview of segmentation process used to segment mice SD-OCT volumes.

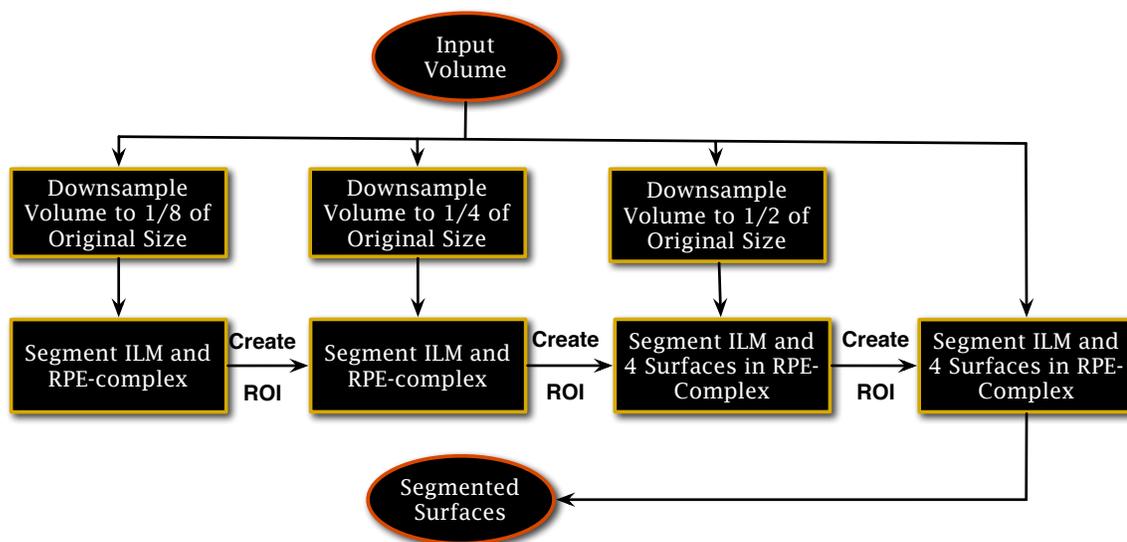


Figure 6.3: Illustration of the segmentation process for the outer retinal surfaces

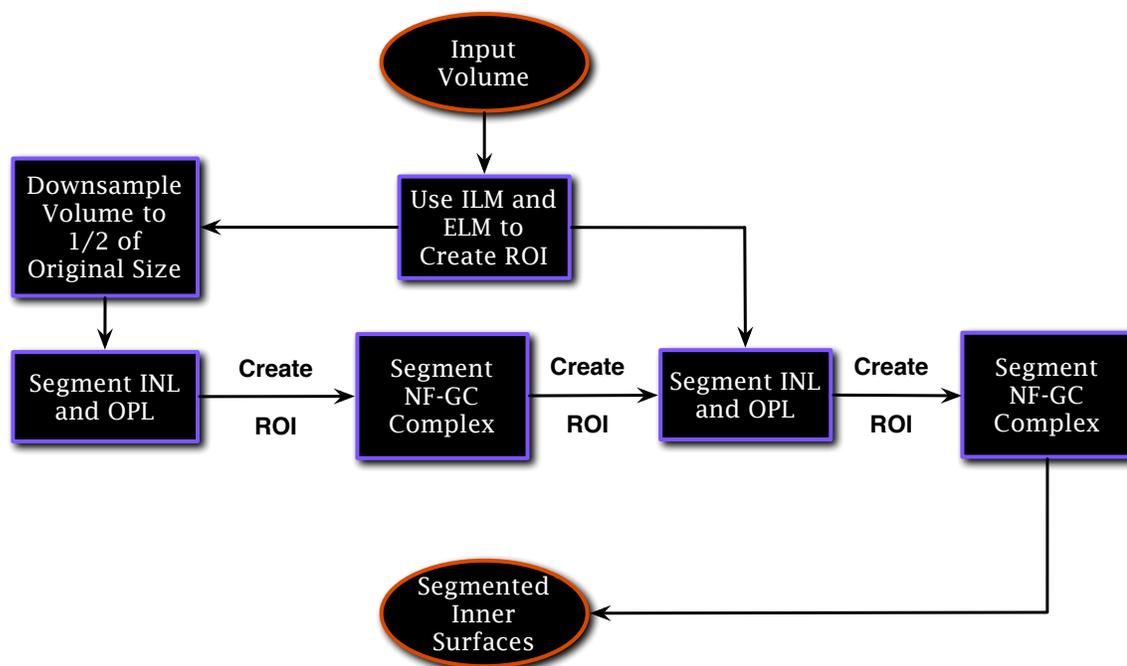


Figure 6.4: Illustration of the segmentation process for the inner retinal surfaces.

simultaneously in the second multi-resolution step (volume down sampled by factor of 4). These four surfaces were segmented simultaneously in the next two stages of the multi-resolution pyramid. Finally, the segmentation of the IS/OS junction was used to create a small region of interest within which the ELM was segmented.

Since the ILM and the ELM are known to encompass the inner retinal layers, these two surfaces were used to create a region of interest (see Fig. 6.4). The inner layers are in general quite thin (with the exception of the IPL) and become indiscernible in volumes down sampled by factors over 2. Therefore, the multi-resolution pyramid used to segment the inner surfaces only contained 2 steps with one volume down sampled by a factor of 2 and the original volume itself. First, the bounding surfaces of the INL and OPL are segmented in the down sampled volume. Since the NF-GC complex lies between the ILM and the interface between the upper boundary of the IPL, these two surfaces can be used to guide the segmentation of the NF-GC

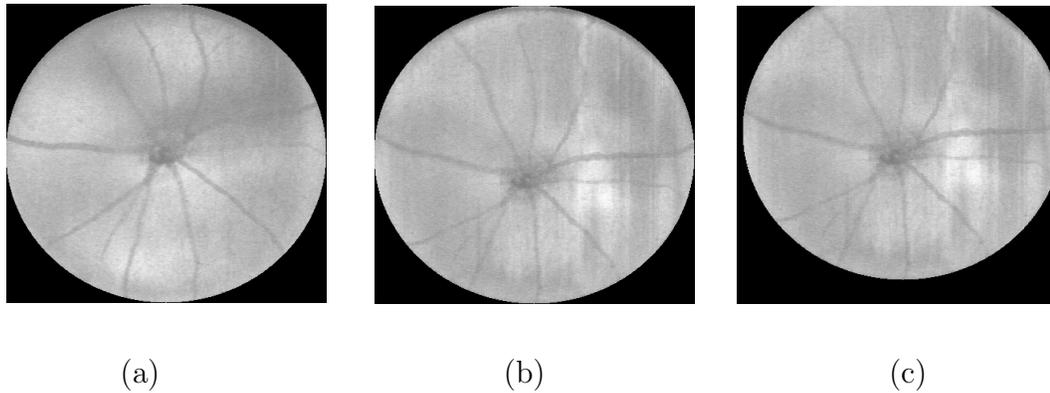


Figure 6.5: Projection images obtained from volumetric SD-OCT images of mice. (a) A correctly centered scan, (b) an off-center scan and (c) corrected scan.

complex. This process was then repeated in the original volume to give us the final segmentations of the NF-GC complex, IPL, INL and OPL.

Some scans were not accurately centered on the ONH during acquisition and needed to be centered manually before the reproducibility could be assessed. For this, the automated segmentation of the IS/OS junction was used to create projection images using a small number of A-scans near this surface. The center of the ONH was manually identified in the projection images and used to translate the volume into the correct location. Fig. 6.5 shows projection images obtained from the repeat scans of the right eye of a mouse, where Fig. 6.5(a) shows one that was acquired correctly while Fig. 6.5(b) shows an example of one that was off-center and needed to be corrected (Fig. 6.5(c)). Finally, correctly centered thickness maps were created for the layers of interest.

6.2 Experimental Methods

SD-OCT images were obtained from mice included in a diabetes study, where type 1 diabetes was induced in 8-10 week old mice C57BL/6J (n=10; Jackson Laboratory, Bar Harbor, ME) by intraperitoneal injection of streptozotocin 200 mg/kg (STZ,

Sigma-Aldrich, St. Louis, MO) dissolved in sodium citrate buffer (pH 4.5). Age-matched C57BL/6J mice (n=8) were assigned to the control group. The volumetric scans centered on the optic nerve head (ONH) (1.4mm x 1.4mm x 1.57mm) were obtained from both eyes (right eye imaged first) of the mice using a SD-OCT scanner (Biotigen Inc., Morrisville, NC, USA) where the scans contained 400 x 400 x 1024 voxels. After the right and then left eye were imaged, the right eye was re-imaged using the same protocol listed above, giving us repeat scans of the right eyes. The diabetic mice were imaged at the onset of the study and 6 weeks after the induction of diabetes, while the normals were imaged once at the onset of the study. The scans obtained from the diabetic mice before induction of diabetes were grouped with the normal scans, providing a total of 18 and 10 repeat scans from the normal and diabetic mice, respectively.

6.2.1 Accuracy Assessment

The accuracy of the segmentation method was assessed by comparing the automated segmentation results to the manual tracings obtained from a retinal specialist. For this, 10 slices were extracted from 10 normal volumetric scans obtained from 10 mice selected at random (from the group of 18). Each of the scans was divided into 10 evenly spaced subvolumes from which one slice was chosen at random. Up to ten intra-retinal surfaces (depicted in Fig. 6.1) were then manually traced on each of these slices. The signed and unsigned border position errors were then computed in microns (μm) for each surface on these slices. The comparisons were limited to an annular region of the retina defined by two circles 0.2mm and 1.2mm in diameter centered on the ONH in order to exclude the ONH region and the most peripheral regions of the scan.

6.2.2 Reproducibility Study

The purpose of the reproducibility study is to quantify the difference in layer thickness measurements obtained on the repeat scans for the normal and diabetic mice. Thickness maps were first created from the centered scans for four retinal layers of interest: 1) the combination of the NF-GCL and IPL, 2) the inner retina (as depicted in Fig. 6.1) which is bounded by the ILM and the bottom of the OPL, 3) the outer retina which extends from the ONL to Bruch's membrane and 4) the total retinal thickness defined from the ILM to Bruch's membrane. The mean thickness was then computed for each of the layers in 8 regions, namely, the central superior (CS), central nasal (CN), central inferior (CI), central temporal (CT), paracentral superior (PS), paracentral nasal (PN), paracentral inferior (PI) and paracentral temporal (PT) regions, as shown in Fig. 6.6. The circles used were 0.2mm, 0.6mm and 1.2mm in diameter, respectively. The reproducibility of the the retinal layer thicknesses were then computed in these 8 pre-defined regions. The statistical parameters computed included the mean difference and standard deviation (expressed in μm) between the thickness measurements obtained from the two repeat scans, and the coefficient of variation (CV), expressed as the ratio of the standard deviation of the difference of the thickness measurements to the mean thickness.

For completeness, these parameters were also computed within the nine retinal layers (created by the 10 retinal surfaces).

6.3 Results

For the accuracy assessment, the automated segmentation results were compared to the manual tracings obtained from the expert on 97 slices. Three slices of the original 100 were excluded, as the expert could not confidently trace all ten surfaces on these slices. Fig. 6.7 shows the manual tracings obtained from the expert and the automated segmentation result on three slices from a normal mouse scan. The mean error for each of the 10 surfaces across all 97 slices is summarized in Table 6.1, and

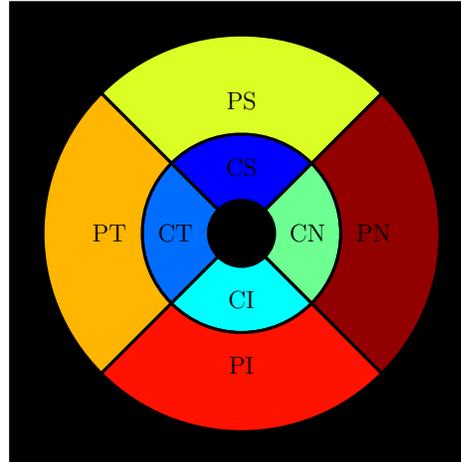


Figure 6.6: Labeled sectoral map within which the reproducibility of thickness measurements were assessed. The reproducibility measures were computed in the 8 regions shown above, namely, the central superior (CS), central nasal (CN), central inferior (CI), central temporal (CT), paracentral superior (PS), paracentral nasal (PN), paracentral inferior (PI) and paracentral temporal (PT) regions. The inner, middle and outer circles are 0.2mm, 0.6mm and 1.2mm in diameter, respectively.

the overall mean unsigned and signed border position error was found to be $3.58 \pm 1.33 \mu\text{m}$ and $-1.00 \pm 1.10 \mu\text{m}$, respectively.

For the reproducibility assessment, one normal mouse was excluded as one of the repeat scans showed a large loss of signal strength, leaving 17 normal and 10 diabetic mice to be employed. The mean thickness and reproducibility statistics of the four retinal layers and the nine retinal layers are summarized in Table 6.2 and Table 6.3, respectively. The mean thickness of the NF-GCL+IPL and total retinal thickness in the normal and diabetic mice were found to be $64.14 \pm 1.17 \mu\text{m}$ and $211.04 \pm 5.89 \mu\text{m}$, and $62.33 \pm 0.99 \mu\text{m}$ and $207.95 \pm 4.16 \mu\text{m}$, respectively. The overall difference in the thickness measurements for four retinal layers is depicted in Fig. 6.8, where the largest thickness difference (under $5 \mu\text{m}$) observed was in the total retinal thickness in the central superior and paracentral inferior regions in the normal and diabetic mice, respectively. The overall mean thickness difference over the 4 retinal layers in

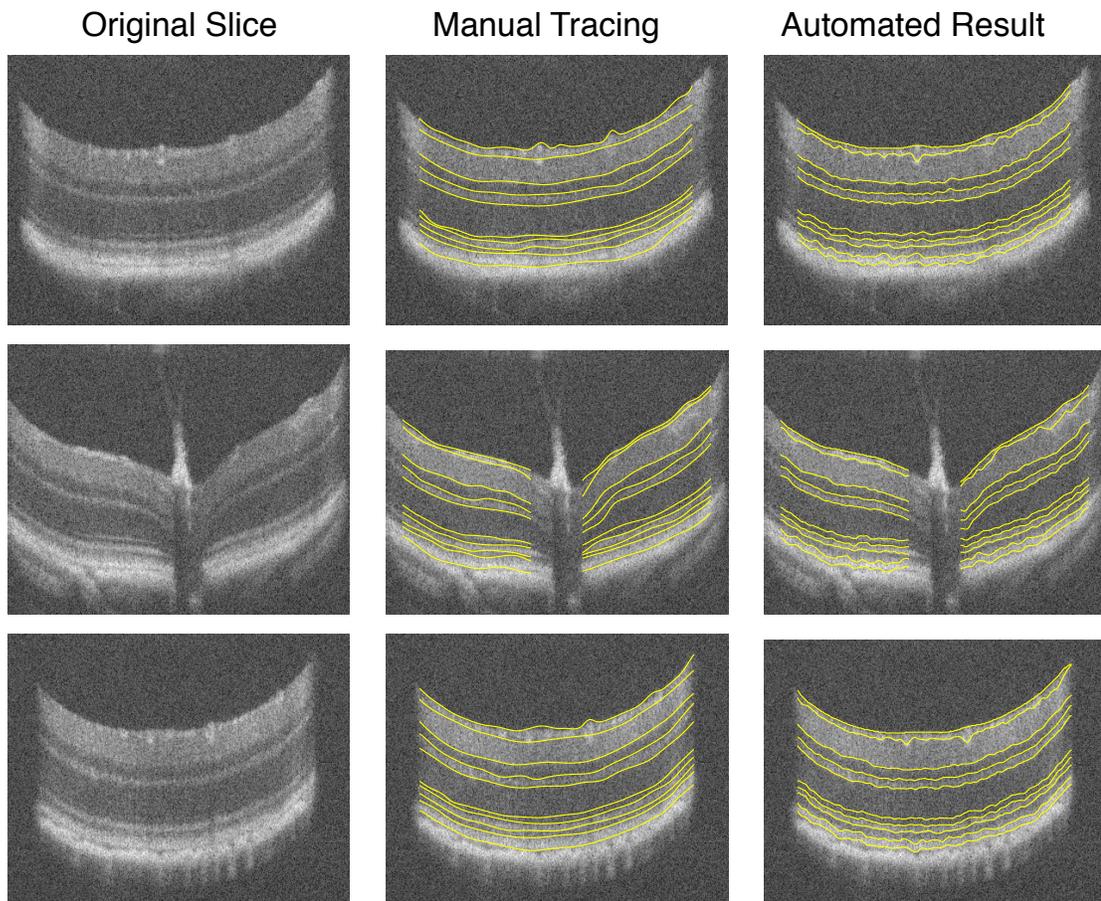


Figure 6.7: Segmentation results of the mice SD-OCT scans. Three slices from a normal mouse SD-OCT image. The original slices (left column) are shown alongside the manual tracings obtained by a retinal specialist (middle column) as well as the automated result (right column).

Table 6.1: Signed and unsigned border position error [†] (mean \pm standard deviation in μm) for each of the 10 segmented intra-retinal surfaces. The errors were computed on a total of 97 slices obtained from 10 normal scans.

	Unsigned Errors (μm)	Signed Errors (μm)
1	2.48 ± 0.72	-1.00 ± 1.10
2	6.97 ± 2.20	5.98 ± 3.32
3	2.86 ± 0.55	-0.42 ± 1.06
4	2.80 ± 0.73	1.09 ± 1.47
5	3.16 ± 0.97	1.62 ± 1.70
6	3.08 ± 1.08	0.78 ± 0.91
7	2.72 ± 0.98	-0.56 ± 0.61
8	3.39 ± 1.20	2.36 ± 0.98
9	4.50 ± 0.97	-2.78 ± 1.92
10	3.87 ± 1.36	0.98 ± 2.34
Average	3.58 ± 1.33	-1.00 ± 1.10

[†] The errors were only computed within an annular region defined by two circles 0.2mm and 1.2mm in diameter.

Table 6.2: Summary of the reproducibility analysis for four retinal layers. The mean thickness (MT), thickness reproducibility (TR) (mean difference \pm SD μm) and the CV (Coefficient of Variation) (%) values computed for the normal (n=17) and diabetic mice (n=10) are shown.

Normal Mice				
	NF-GCL+IPL	Inner Retina	Outer Retina	TRT
MT	64.14 ± 1.17	101.02 ± 1.57	110.00 ± 4.87	211.04 ± 5.89
TR	0.99 ± 0.25	1.17 ± 0.31	2.33 ± 0.74	2.98 ± 0.75
CV	1.49%	1.04 %	1.89%	1.18%
Diabetic Mice				
	NF-GCL+IPL	Inner Retina	Outer Retina	TRT
MT	62.33 ± 0.99	98.78 ± 1.55	109.17 ± 3.19	207.95 ± 4.16
TR	1.27 ± 0.20	1.73 ± 0.33	2.33 ± 0.88	3.27 ± 0.90
CV	1.79%	1.77 %	3.06%	1.56%

Table 6.3: Summary of the reproducibility analysis for 11 retinal layers. The mean thickness (MT), thickness reproducibility (TR) (mean difference \pm SD μm) and the CV (Coefficient of Variation) (%) values computed for the normal (n=17) and diabetic mice (n=10) are shown.

Layer	Normal Mice			Diabetic Mice		
	MT	TR	CV	MT	TR	CV
1	11.71 \pm 2.64	2.28 \pm 1.87	25.19%	10.79 \pm 3.66	2.97 \pm 2.36	21.86%
2	52.48 \pm 2.83	2.02 \pm 1.71	4.87%	51.89 \pm 4.08	3.38 \pm 2.41	4.64%
3	20.15 \pm 1.25	0.44 \pm 0.40	2.97%	19.39 \pm 0.85	0.43 \pm 0.47	2.45%
4	16.77 \pm 0.65	0.40 \pm 0.53	3.97%	17.18 \pm 0.71	0.48 \pm 0.64	3.73%
5	52.31 \pm 4.87	1.80 \pm 1.90	4.93%	53.07 \pm 4.33	2.53 \pm 2.99	5.62%
6	12.37 \pm 0.12	0.10 \pm 0.12	1.08%	12.51 \pm 0.22	0.22 \pm 0.16	1.30%
7	10.72 \pm 0.52	0.45 \pm 0.39	4.79%	10.52 \pm 0.66	0.41 \pm 0.41	3.90%
8	21.92 \pm 2.83	1.29 \pm 1.23	8.03%	21.14 \pm 2.10	1.28 \pm 1.07	5.06%
9	12.65 \pm 1.25	0.36 \pm 0.29	3.53%	12.54 \pm 0.75	0.48 \pm 0.38	3.04%

the normal and diabetic mice was found to be $1.86 \pm 0.95 \mu\text{m}$ and $2.15 \pm 0.86 \mu\text{m}$, respectively.

The region-based CV values for the normal and diabetic mice are depicted in Fig. 6.9. The CVs for all the layers, with the exception of the NF-GC complex, were under 10% with the highest values being noted among the diabetic mice. The CV for the NF-GCL+IPL and total retinal thicknesses was found to be 1.49% and 1.18%, and 1.79% and 1.56% in the normal and diabetic mice, respectively.

6.4 Discussion

The adapted graph-theoretic approach described here for the segmentation of intra-retinal layers in mouse SD-OCT volumetric scans proved to be quite accurate and reliable. The reproducibility of the thickness measurements is comparable with the voxel size ($1.57 \mu\text{m}$) and the CV values are under 5%, with the exception of the NF-GC complex. The accuracy assessment also shows that the NF-GC complex is the most difficult to segment, with a signed error indicating that the algorithm tends to overestimate this layer. This is not unexpected as the layer is quite thin, disappearing

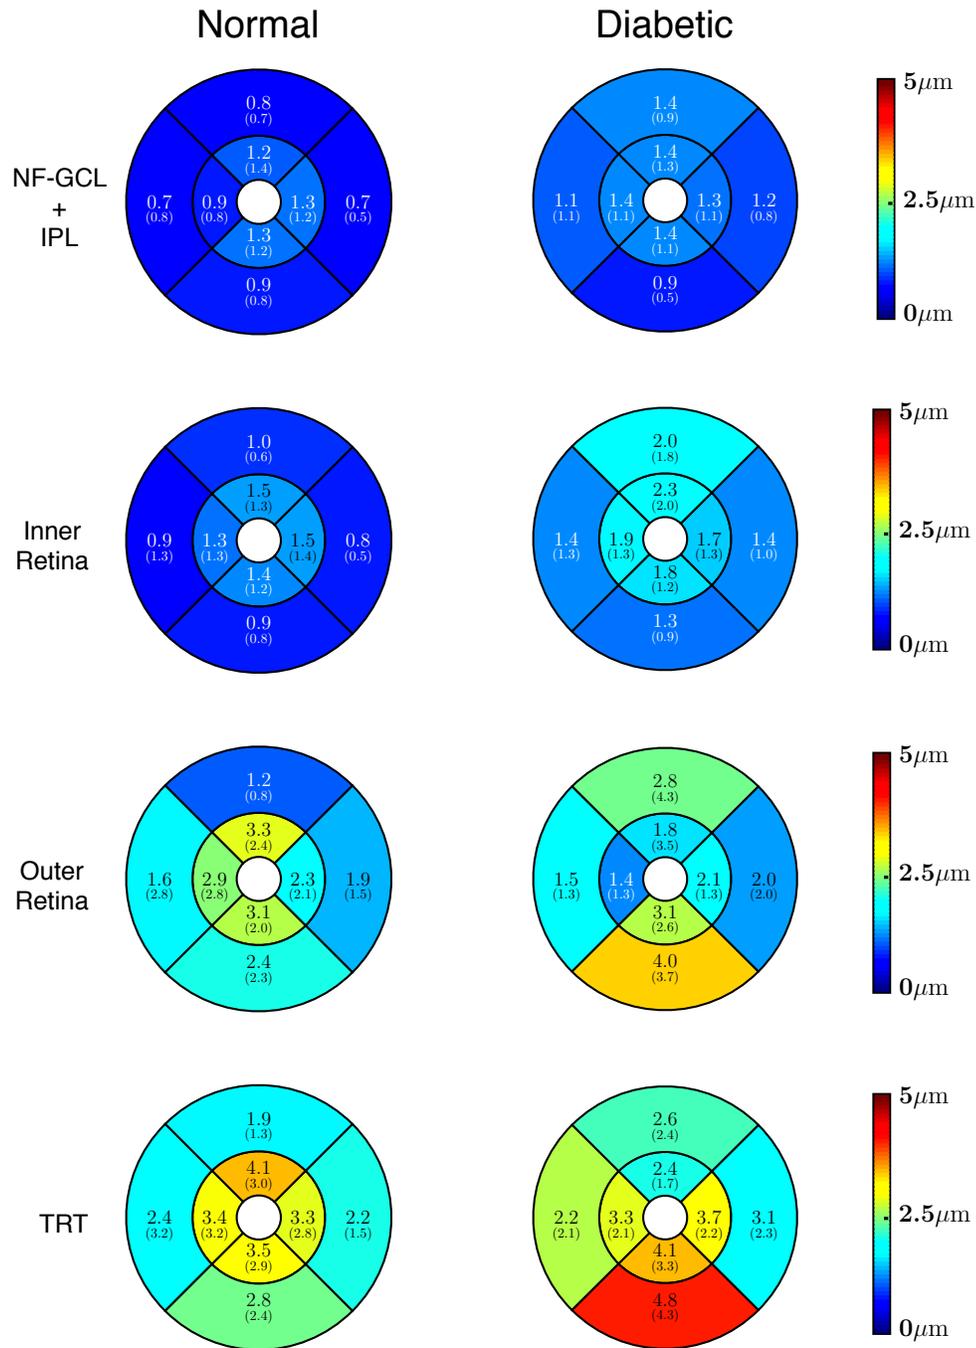


Figure 6.8: The thickness differences computed in four retinal layers. The differences were computed (mean (SD) μm) within the NF-GCL+IPL (top row), the inner retina (second row), the outer retina (third row) and the total retinal thickness (last row) for the normal (left) and the diabetic (right) mice.

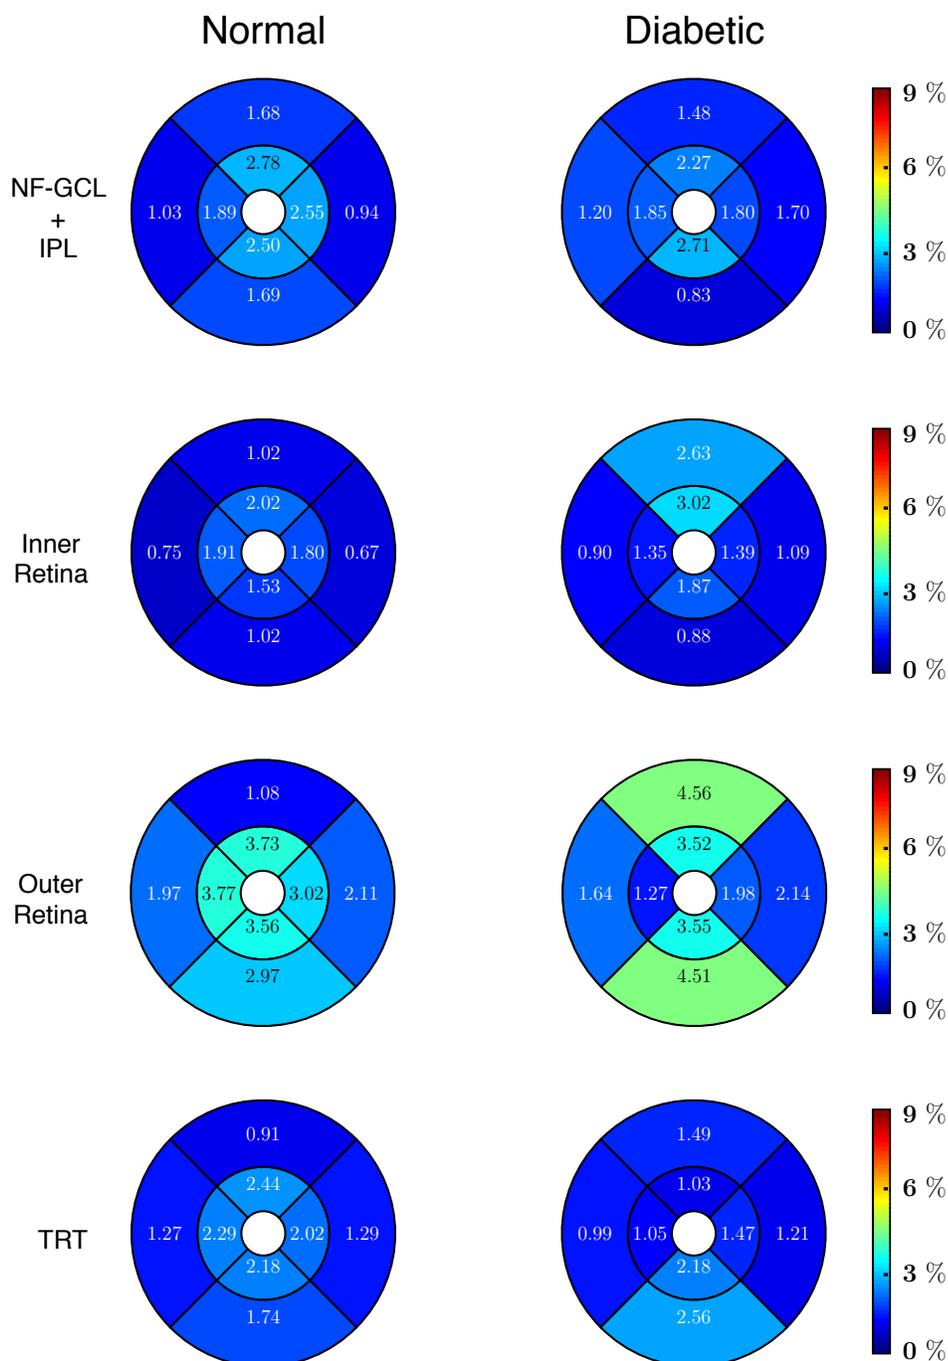


Figure 6.9: The coefficient of variation (CV%) computed in four retinal layers. The values were computed within the NF-GCL+IPL (top row), the inner retina (second row), the outer retina (third row) and the total retinal thickness (last row) for the normal (left) and the diabetic (right) mice. The CV values were obtained by dividing the standard deviation of the thickness difference per region by the mean thickness of the region.

completely at the periphery of the scans and can even be quite difficult to visually discern in some locations. The incorporation of learned information, such as textural features and expected surface topology (described in Chapters 7 and 8), however, does reduce these errors. The reproducibility of the thickness measurements showed slightly larger errors in the diabetic mice. This is also not unexpected as the presence of disease does make it difficult to segment the layers, and it is also important to note that the total number of scans obtained in diseased mice were lower than the total number obtained from normals. Nevertheless, the reproducibility is expected to be sufficient for the monitoring of pathophysiologically meaningful change over time [90].

CHAPTER 7

A COMBINED MACHINE-LEARNING AND GRAPH-THEORETIC APPROACH FOR THE SEGMENTATION OF RETINAL SURFACES IN SD-OCT VOLUMES

As described in the previous chapter, the graph-theoretic approach is an effective method that can be used to segment retinal surfaces in SD-OCT images of humans as well as mice. However, the adaptation of this method to new models is not a straight forward process and can be quite challenging. In addition to the different anatomy which needs to be taken into consideration when determining the feasibility constraints, designing the cost function can be an arduous and time-consuming task. For instance, the cost function has been designed using Sobel filters coupled with a de-noising method [21,27] when segmenting human scans. However, in the mice scans, Gaussian derivatives proved to be a better choice. This task is further confounded in the presence of disease or when the structures of interest are difficult to identify, such as the NF-GC complex, which not only showed the largest unsigned border position error but also had the lowest thickness reproducibility. Utilizing a machine-learning based approach to design the cost function can not only have a significant impact on the segmentation accuracy (as shown in Chapter 5), it provides the added advantage of being quickly adapted to new models as the only requirement is the training set.

This chapter details a combined machine-learning and graph-theoretic method for the automated segmentation of retinal surfaces in SD-OCT images, where a machine-learning based approach is used to design the in-region and on-surface cost terms of the cost function. In addition to designing the cost terms, it is also important to gauge the impact of these learned features on the final segmentation accuracy to prevent the inclusion of terms that do not significantly improve the accuracy or in some cases, hurt the segmentation accuracy. This allows for the graph-theoretic method to be adapted and tailored to new models by retraining the system on a new dataset and finding the best set of learned components that provide the best

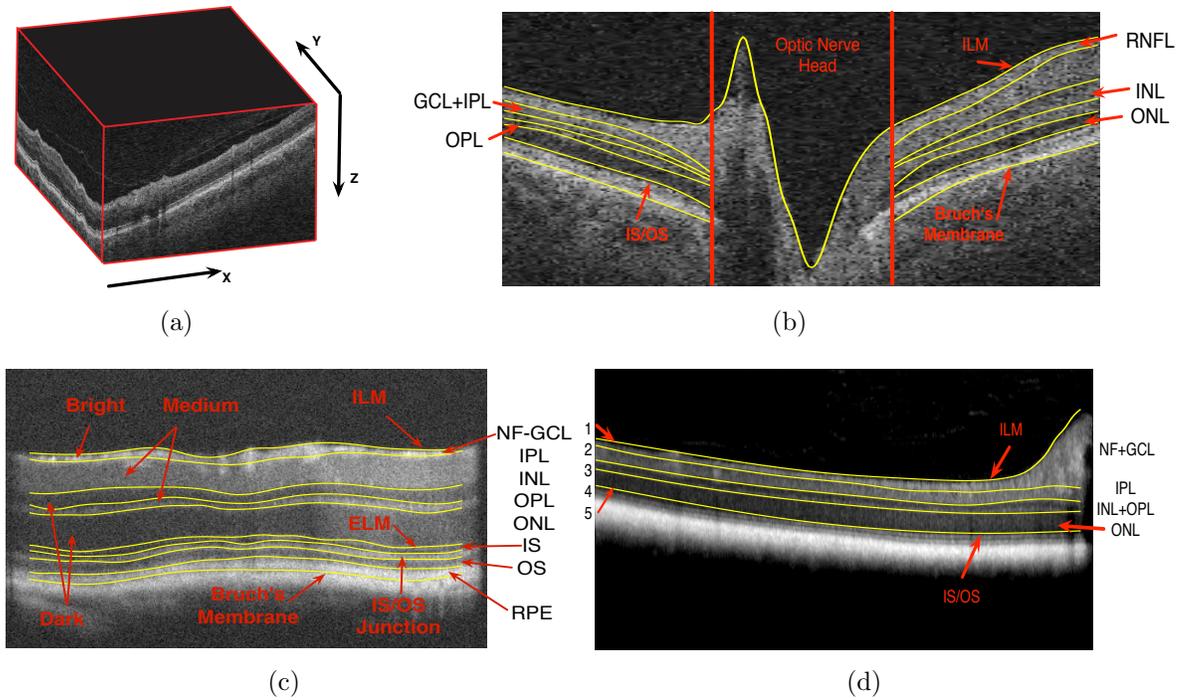


Figure 7.1: Examples of images used to evaluate the efficacy of the proposed method. (a) Illustration of a 3-D volumetric SD-OCT scan of a human eye centered on the optic nerve head. A central slice from a volumetric scan obtained from (b) a human subject showing 7 retinal surfaces, (c) a mouse showing 10 retinal surfaces and (d) a canine showing 5 retinal surfaces.

segmentation accuracy. The adaptability of the method is demonstrated using three sets of images (Fig. 7.1 shows examples of each) obtained from: 1) human patients that presented with glaucoma, 2) mice and 3) canines employed in a glaucoma study. The overall error obtained using the recommended framework of the graph-theoretic method on these three sets was found to be $6.45 \pm 1.87 \mu\text{m}$, $3.35 \pm 0.62 \mu\text{m}$ and $9.75 \pm 3.18 \mu\text{m}$, respectively.

7.1 Methods

The features used to train the classifier included Haar, Gabor, steerable-Gaussian and intensity-base features as described previously in Section 5.1. The segmentation results from Chapter 5 indicated that the grouping of regions based on similar prop-

erties, such as intensity, helps reduce misclassifications as well significantly increases the segmentation accuracy. Thus, the regions were again grouped on the basis of intensity into high-, medium- and low-intensity regions. Similarly, as surfaces lie either below a high-intensity region with a lower intensity region below it, or vice versa, they were categorized as either dark-to-bright interfaces or bright-to-dark interfaces.

k -NN classifiers are simple classifiers but can prove to be quite efficient in a number of applications. Here, the classes showed large size disparities and the potentially large dimensionality of the data led us to consider other classification models. Thus, we utilized a random forest [56], which is an ensemble classifier composed of a number of decision trees. The class label of an input vector is determined by the majority vote cast by the individual trees, and regression can be performed by computing the ratio of the majority votes to the total number of trees in the forest. We utilized the OpenCV implementation of random trees. The random forest was initially trained using all the texture features corresponding to the three region categories. The individual feature significances were used to limit the feature set to those that contributed at least 2% to the overall feature significance. This final set of features were then used to retrain the random forest and create probability maps for the high-, medium- and low-intensity regions. A similar process was used to create probability maps for the dark-to-bright and bright-to-dark interfaces. Fig. 7.2 show the probability maps obtained for scans acquired from a human, a mouse and a canine.

7.1.1 Segmentation Process

The retinal surfaces segmented in the human, mice and canine scans are depicted and labeled in Fig. 7.1(b), (c) and (d), respectively. Although the number of target surfaces differ for the different datasets, the overall segmentation process used a two-stage approach [21,27] (see Section 6.1), which begins by segmenting the outer surfaces, namely the internal limiting membrane (ILM), the junction of the inner and outer segments (IS/OS junction) and the bounding surfaces of the retinal pigment

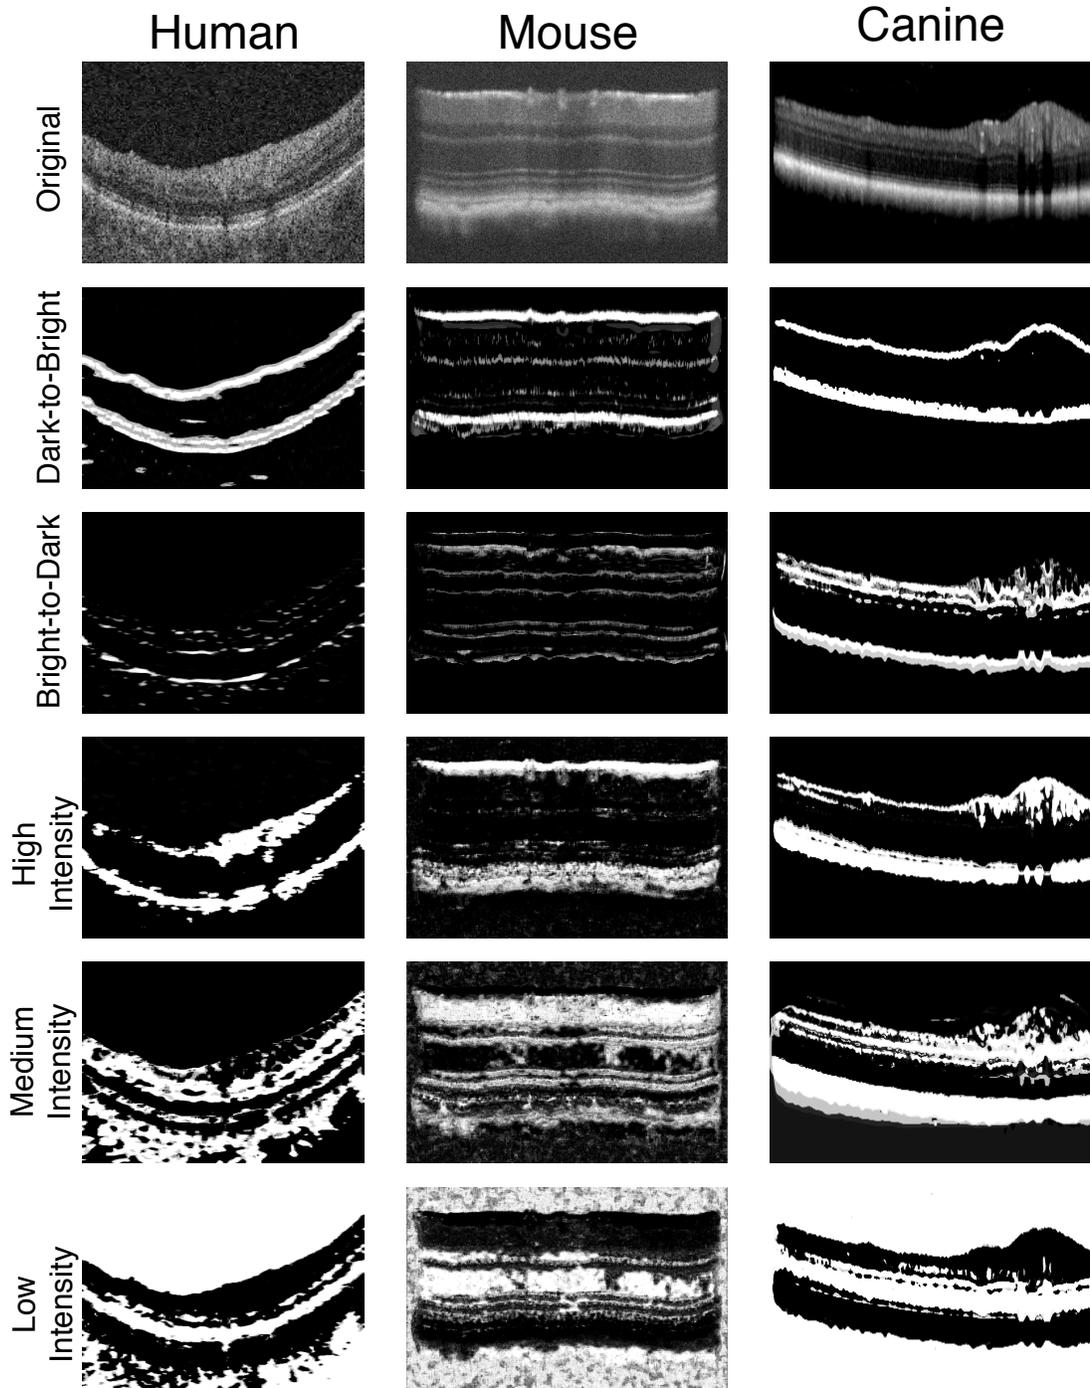


Figure 7.2: Probability maps obtained for the three categories of regions and two categories of surfaces for scans acquired from a human, a mouse and a canine. The top row shows slices from SD-OCT volumes obtained from (left) a human, (middle) a mouse and (right) a canine. The probability maps obtained for the dark-to-bright surfaces and bright-to-dark are shown in the second and third rows, respectively. The probability maps obtained for the high-, medium and low-intensity regions are as shown in the fourth, fifth and sixth rows, respectively.

epithelium (RPE), followed by the segmentation of the inner layers, namely the retinal nerve fiber layer (RNFL), the inner plexiform layer (IPL), the inner nuclear layer (INL) and the outer plexiform layer (OPL).

As the surfaces that mark the bounds of the inner and outer segments of the photoreceptors, and the RPE are extremely close together, learning regional features was difficult and the segmentation of these surfaces could not incorporate a learned in-region cost term. The inner layers however, are ideal for learning in-region as well as on-surface cost terms and the segmentation of these surfaces included either a learned in-region or an on-surface cost term or a combination of both.

A cross-validation approach was utilized for the experiments, where n datasets were utilized to train the classifier and the $N - n$ datasets were used to test the system. In the case of the human and mice data, $n = 2$, while $n = 5$ for the dog datasets (see Section 7.2.1 for further details on the data used). The automated feature selection based on the individual feature significance obtained in each of these experiments yielded a total of 20 and 19 features for the in-region and on-surface cost terms, respectively, in the human dataset. Similarly, for the mice data, 21 features were used to learn the in-region and on-surface terms, while in the canine datasets, 17 features were used to learn these cost terms.

In addition to finding a good set of features that describe the regions and surfaces, it is also important to determine the optimal value for α in Eq. 3.3, that defines the combination of the in-region and on-surface cost terms [21]. As each of the surfaces may prefer a different on-surface to in-region cost term ratio, it may be preferable to find the optimal set of weights β_i for each of the surfaces. The cost function can now be expressed as:

$$C_T = \sum_{i=1}^n \beta_i C_{S_i} + \sum_{j=0}^n C_{R_j}, \quad \text{and} \quad \beta_i = \begin{cases} \frac{\alpha_i}{(1-\alpha_i)}, & \text{if } 1 - \alpha_i > 0 \\ C_{large}, & \text{if } 1 - \alpha_i \approx 0. \end{cases} \quad (7.1)$$

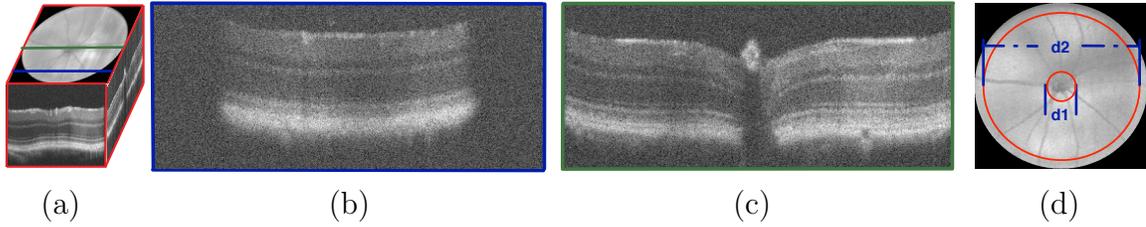


Figure 7.3: The evaluation criteria for scans of mice. (a) Illustration of a volumetric SD-OCT scan obtained from a mouse, showing the circular “valid” region of the projection image. The circular nature of the scan results in large noisy regions in slices near the (b) the periphery. The optic nerve head causes shadows in the (c) central slices of the scan. (d) The evaluation of the segmentation accuracy was therefore, limited to an annular region defined by $d1=0.2\text{mm}$ and $d2=1.2\text{mm}$ in order to avoid the ONH and the noisy peripheral regions.

where α_i is the value that provides the lowest segmentation error for surface i . For this, we used Eq. 3.3 and tested a range of α in increments of 0.1. When an estimate of α was found, smaller steps were used to further refine the optimization.

7.2 Experimental Methods

7.2.1 Data Description

The proposed method was evaluated on three sets of images obtained from human patients, mice, and canines as detailed below:

The first set consisted of 10 optic nerve head (ONH) OCT volumes acquired from 10 human glaucoma patients on a spectral-domain Cirrus (Carl Zeiss Meditec, Inc., Dublin CA) scanner [27]. This set was used in our previous work [27, 81] and is detailed in Section 5.2.

The second set consisted of 10 ONH-centered OCT volumes acquired from 10 normal mice. This dataset was used in the evaluation of the baseline approach and is described in Section 6.2. Ten retinal surfaces, as depicted in Fig. 7.1(c), were segmented on these scans; five within the inner retina, namely the bounding surfaces

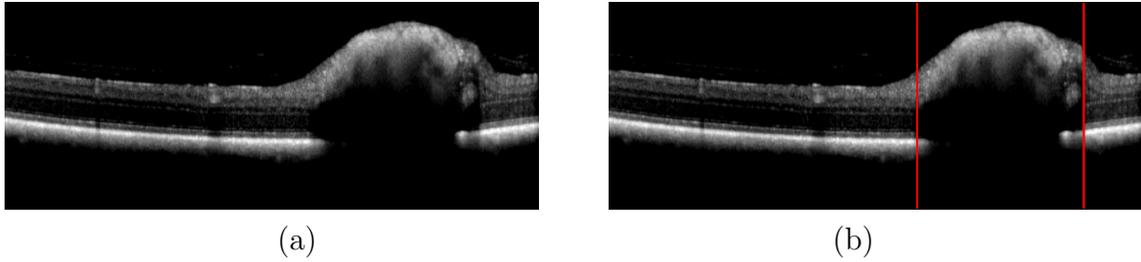


Figure 7.4: The segmentation accuracy evaluation criteria for scans of canines. Regions of low signal strength and the optic nerve head region were excluded from the validation. (a) A slice from an SD-OCT scan obtained from a canine, showing the optic nerve head region within which surfaces are difficult to discern and the (b) corresponding mask. The region within the red lines was excluded.

of the nerve fiber-ganglion cell complex (NF+GCL), the inner plexiform layer (IPL), the INL, and the OPL; and five in the outer retina, namely, the external limiting membrane (ELM), the IS/OS junction, the bottom of the outer segments and the upper and lower boundaries of the RPE. Due to the curvature of the retina, only a circular region within the 1.4mm x 1.4mm scanned area contains valid information. Fig. 7.3(a) depicts a volumetric scan and the circular projection image usually seen in these images. Figs. 7.3(b) and (c) show examples of a peripheral slice and a central slice, respectively. Thus, the evaluation of the segmentation accuracy was limited to an annular region (see Fig. 7.3(d)) defined by two circles 0.2mm and 1.2mm in diameter in order to exclude the ONH and the noisy peripheral regions. The segmentation accuracy was assessed on 10 volumetric scans, where manual tracings were obtained from a retinal specialist on a total of 100 slices. For each of the datasets, 10 slices were chosen at random from sections 30 slices deep while avoiding the first and last 50 slices as the section of the retina imaged in the first and last 50 slices comprise less than 50% of the B-scan.

The third data set consisted of SD-OCT retinal scans obtained from a colony of nineteen Basset Hounds. The colony consisted of three nonaffected canines, six

canines with genetic susceptibility to glaucoma that did not develop the disease and 10 canines with genetic susceptibility that presented with glaucoma. Baseline scans of each animal were obtained at 6 months of age. Subsequent SD-OCT scans were obtained after provocative testing by acute elevation of intraocular pressure (AEIOP). The scans were centered on the macula and were obtained on a Heidelberg SD-OCT Spectralis scanner (Heidelberg Engineering, Germany). Varying protocols were used to acquire these scans, but the typical scan encompassed a region 4.2mm x 4.5mm x 1.9mm and contained 768x19x496 voxels. An independent observer (experienced in retinal anatomy of canines) traced five surfaces (as shown in Fig. 7.1(d)), namely the bounding surfaces of the NF+GCL, the IPL and the ONL. The INL and the OPL were also segmented but they are extremely thin and the observer was unable to trace these surfaces consistently. Thus, the accuracy evaluation was limited to five surfaces. Furthermore, certain sections of the images showed low signal strength as depicted in Fig. 7.4. These regions where the surfaces were not visible were excluded from the evaluation as well.

7.2.2 Accuracy Evaluation

A cross-validation scheme was used to train the classifiers and evaluate the final segmentation accuracy. For the human and mice datasets that contain ten images each, eight scans were used to train the system and the remaining two were segmented after the incorporation of the learned cost terms. This was repeated four times until all the scans were segmented. In the case of the canine scans, 14 datasets of the 19 were used to train the classifiers and the remaining datasets were segmented. This was repeated three more times to provide segmentation results on all of the 19 datasets.

The unsigned border position errors between the automated segmentations and the manual tracings were used as the measure of the segmentation accuracy. As the cost function can be represented as a weighted combination of on-surface and in-region cost terms, in addition to evaluating the overall segmentation accuracy,

it is also important to gauge the individual impact of these automatically designed terms. This will help not only help to tailor the graph-theoretic method to specific applications, it will also prevent the inclusion of features that may potentially hurt the segmentation results. This was done by comparing the segmentation results obtained when using the new cost function designs to those obtained when using a baseline approach. Note that the baseline approaches did not incorporate any learned cost terms but utilized either an on-surface cost term [91] or a weighted combination of an in-region and an on-surface cost term [27]. Three possible frameworks were, therefore tested, where the graph-theoretic method incorporated:

1. a combination of learned in-region cost terms and Gaussian-derivative based on-surface terms (Experiment I),
2. only learned on-surface cost terms (Experiment II), and
3. a combination of learned in-region and on-surface cost terms (Experiment III).

Baseline approaches have been previously described for the human [27] and mice scans (Section 6), however, designing one for the canine dataset proved to be extremely difficult. Thus, this dataset does not have a baseline approach. Paired *t*-tests that were corrected for multiple comparisons were used to compare the results obtained.

7.3 Results

7.3.1 Data Obtained from Human Subjects

Table 7.1 shows the overall border position errors for the 7 surfaces obtained using Experiments I, II and III as well as the results obtained when no learned component was incorporated into the segmentation method [27] (labeled ‘Baseline’). The overall border position error obtained for Experiments I, II and III was $6.45 \pm 1.87 \mu\text{m}$, $8.26 \pm 2.72 \mu\text{m}$ and $8.67 \pm 2.88 \mu\text{m}$, respectively. The errors observed for the inner surfaces

Table 7.1: Unsigned border position error (mean \pm SD) in μm computed on the human dataset.

Surface	Baseline	Experiment I *	Experiment II °	Experiment III •
1	4.90 ± 1.54	4.90 ± 1.54	4.90 ± 1.54	4.90 ± 1.54
2	14.43 ± 5.63	9.82 ± 3.37	13.28 ± 4.55	12.68 ± 5.51
3	10.96 ± 4.06	6.47 ± 1.99	9.07 ± 2.94	10.96 ± 4.59
4	10.46 ± 2.79	6.42 ± 2.12	9.13 ± 2.82	10.44 ± 4.05
5	10.73 ± 2.78	6.41 ± 2.11	8.55 ± 2.58	8.84 ± 3.44
6	3.87 ± 1.32	3.87 ± 1.32	5.66 ± 0.54	5.66 ± 0.54
7	7.24 ± 1.74	7.24 ± 1.74	7.22 ± 1.32	7.22 ± 1.32
Mean	8.94 ± 3.76	6.45 ± 1.87	8.26 ± 2.72	8.67 ± 2.88

* The cost function incorporated a combination of learned in-region cost terms and Gaussian-derivative based on-surface terms.

° The cost function incorporated the learned on-surface cost term.

• The cost function consisted of a weighted combination of the learned in-region and on-surface cost term.

for all three experiments were also comparable with the inter-observer variability [27] that was $12.79 \pm 3.36 \mu\text{m}$, $13.74 \pm 2.04 \mu\text{m}$, $9.28 \pm 3.00 \mu\text{m}$ and $7.67 \pm 1.69 \mu\text{m}$ for surfaces 2, 3, 4 and 5, respectively. The segmentation errors observed in Experiments I showed significant improvements ($p < 0.001$) over the baseline results for all four inner surfaces. The incorporation of only the learned on-surface cost results in smaller errors than the baseline results, however, they were not significantly different from the baseline results. In Experiment III, surface 2 (bottom of the RNFL) and surface 5 (top of the ONL) significantly improved in comparison with baseline results. The best results for all surfaces, however, were obtained using the learned in-region cost term.

Figs. 7.5(a) shows a near-central slice from a scan obtained from a human glaucoma patient. Errors were commonly seen near the ONH-region when no learned cost terms were incorporated into the graph-theoretic method [27] as shown in Fig. 7.5(c)

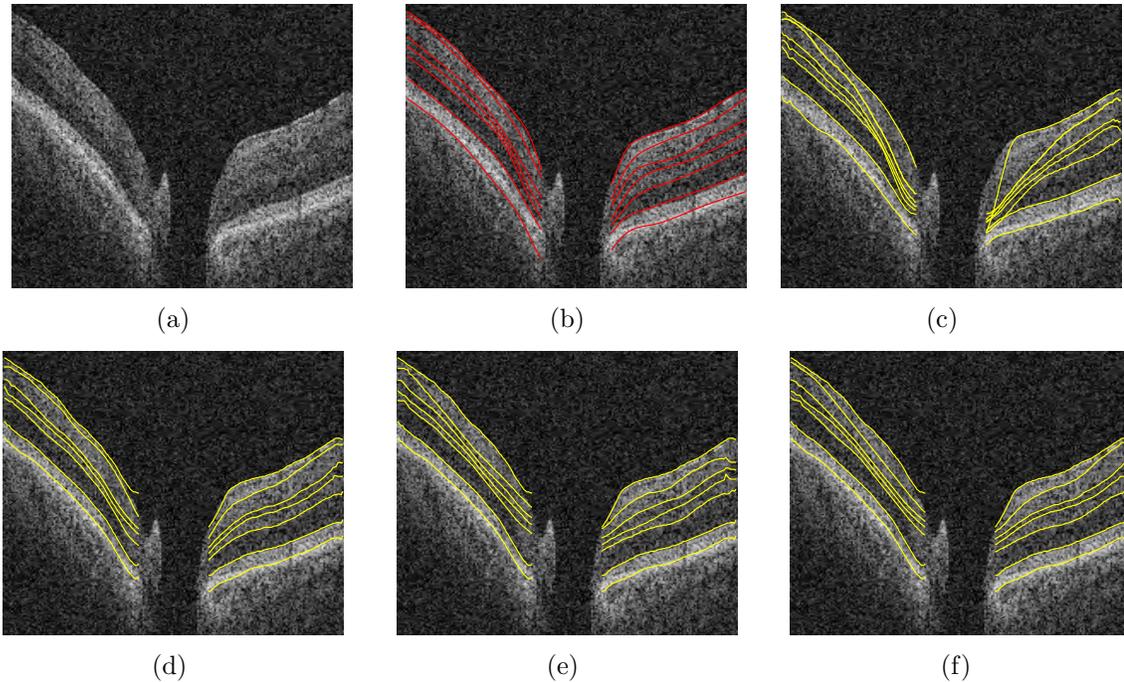


Figure 7.5: Segmentation results on the human dataset. (a) A near-central slice from an SD-OCT scan acquired from a human subject that presented with glaucoma alongside (b) the manual tracings obtained from an independent observer and the automated results obtained using (c) the previously described method [27] that did not incorporate any learned cost terms, (d) learned in-region cost function, (e) learned on-surface cost function and (f) a combination of both.

as the gradients that mark the layer boundaries are quite small and are in fact, comparable to the speckle noise seen in these images. The incorporation of the learned in-region cost term had the most significant impact on the results as shown in Fig. 7.5(d). The incorporation of the learned on-surface cost term and the combination of the in-region and on-surface cost term also improved the segmentation accuracy from the baseline results as depicted in Figs. 7.5(e) and (f), respectively.

7.3.2 Data Obtained from Mice

The unsigned border position errors were computed on a total of 97 slices as not all of the 10 surfaces could be confidently traced on 3 slices of the original 100. The overall border position errors for the 10 surfaces obtained using Experiments I, II and III

Table 7.2: Unsigned border position error (mean \pm SD) in μm obtained on the mice data.

Surface	Baseline	Experiment I [*]	Experiment II [°]	Experiment III [•]
1	2.48 \pm 0.72	2.48 \pm 0.72	2.48 \pm 0.72	2.48 \pm 0.72
2	6.97 \pm 2.20	4.80 \pm 0.86	5.16 \pm 1.34	4.45 \pm 0.56
3	2.86 \pm 0.55	3.97 \pm 0.50	3.30 \pm 0.73	3.24 \pm 0.83
4	2.80 \pm 0.73	2.95 \pm 0.83	2.88 \pm 0.75	2.59 \pm 0.71
5	3.16 \pm 0.97	2.74 \pm 0.64	4.05 \pm 1.18	3.98 \pm 1.33
6	3.08 \pm 1.08	3.08 \pm 1.08	3.01 \pm 0.86	3.01 \pm 0.86
7	2.72 \pm 0.98	2.72 \pm 0.98	2.89 \pm 0.80	2.89 \pm 0.80
8	3.39 \pm 1.20	3.39 \pm 1.20	3.57 \pm 0.98	3.57 \pm 0.98
9	4.50 \pm 0.97	4.50 \pm 0.97	3.58 \pm 1.17	3.58 \pm 1.17
10	3.87 \pm 1.36	3.87 \pm 1.36	3.72 \pm 1.15	3.72 \pm 1.15
Mean	3.58 \pm 1.33	3.45 \pm 0.80	3.46 \pm 0.76	3.35 \pm 0.62

^{*} The cost function incorporated a combination of learned in-region cost terms and Gaussian-derivative based on-surface terms.

[°] The cost function incorporated the learned on-surface cost term.

[•] The cost function consisted of a weighted combination of the learned in-region and on-surface cost term.

were $3.45 \pm 0.80 \mu\text{m}$, $3.46 \pm 0.76 \mu\text{m}$ and $3.35 \pm 0.62 \mu\text{m}$, respectively (see Table 7.2). The border position errors obtained when no learned components were incorporated into the graph-theoretic method [91] was $3.58 \pm 1.33 \mu\text{m}$. The incorporation of the learned cost terms (in Experiments I, II and III) significantly improved ($p < 0.001$) the results for surface 2 over the baseline results. The incorporation of the learned in-region cost term also significantly reduced ($p < 0.001$) the errors noted in surface 5 over the baseline results.

Figs. 7.6(a) and (b) show a slice from a mouse scan and the manual tracings obtained from a retinal specialist. Figs. 7.6(d)-(f) show the results obtained using the learned on-surface cost, in-region cost and a combination of both, respectively. This slice showed higher errors than the average when no learned cost term was incorporated into the graph-theoretic method. The slice is very close to the ONH

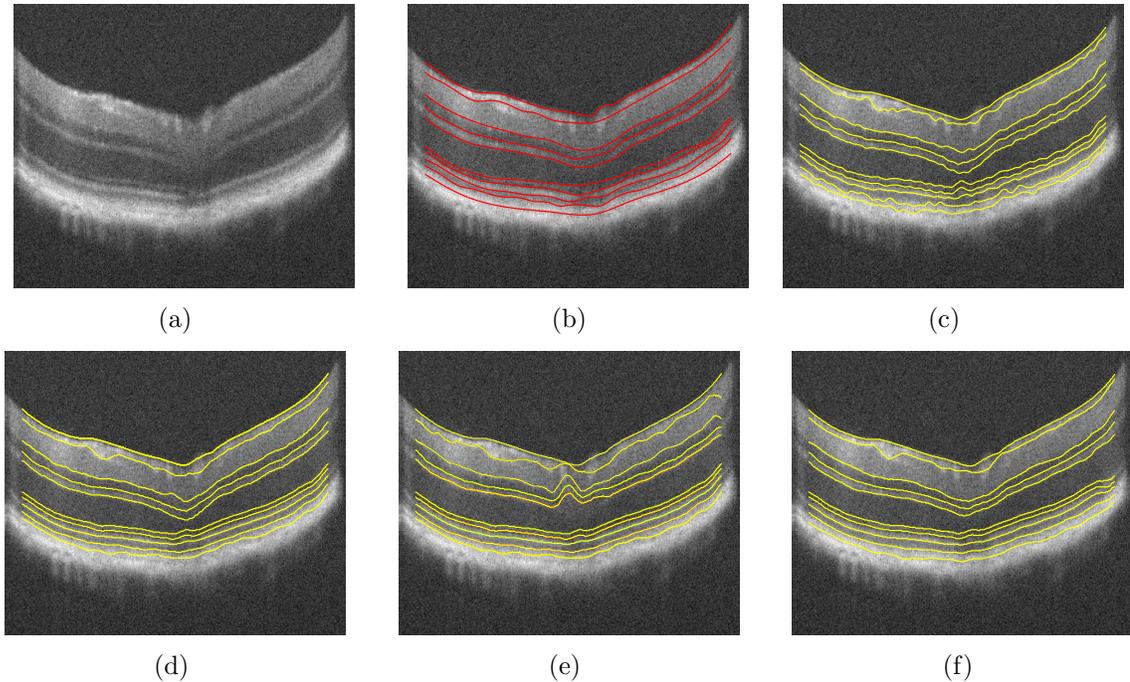


Figure 7.6: Segmentation results obtained for the mice dataset. (a) A near-central slice from an SD-OCT volume of a mouse retina alongside (b) the manual tracings obtained from a retinal specialist and the automated results obtained using (c) an approach that did not incorporate any learned cost terms and after the inclusion of learned (d) in-region cost function, (e) on-surface cost function and (f) a combination of both.

region (0.13mm away from the center of the scan), and the large vessel shadows make the central area difficult to segment. However, the use of the learned in-region and on-surface cost terms corrected these segmentation errors.

7.3.3 Data Obtained from Canines

The overall border position errors for the 10 surfaces obtained using Experiments I, II and III were 9.75 ± 3.18 , 11.39 ± 4.82 and 10.14 ± 4.26 , respectively. The unsigned border position errors for the five surfaces are tabulated in Table 7.3. The segmentation errors noted in surface 2 (bottom of NF+GCL) in Experiment I was significantly smaller ($p < 0.01$) than those noted in Experiments II and III. The errors noted in surface 3 (bottom of the IPL) was smallest when both learned cost

Table 7.3: Overall unsigned border position error (mean \pm SD) in microns obtained on the 19 canine scans.

Surface	Experiment I *	Experiment II °	Experiment III •
1	6.65 \pm 1.88	6.65 \pm 1.88	6.65 \pm 1.88
2	14.50 \pm 3.80	18.05 \pm 3.10	17.31 \pm 4.83
3	10.12 \pm 4.98	14.30 \pm 5.90	9.69 \pm 5.05
4	10.49 \pm 5.30	10.74 \pm 5.47	9.88 \pm 5.98
5	7.00 \pm 2.03	7.00 \pm 2.03	7.00 \pm 2.03
Mean	9.75 \pm 3.18	11.39 \pm 4.82	10.14 \pm 4.26

* The cost function incorporated a combination of learned in-region cost terms and Gaussian-derivative based on-surface terms.

° The cost function incorporated the learned on-surface cost term.

• The cost function consisted of a weighted combination of the learned in-region and on-surface cost term.

terms were incorporated into the graph-theoretic method and was significantly better ($p < 0.001$) than when only learned on-surface costs were incorporated, but did not significantly differ from the errors noted when only the learned in-region cost term was used. The errors noted in surface 4 (upper surface of the ONL) in all three experiments did not significantly differ from each other.

Fig. 7.7 (a) shows a slice from an SD-OCT volumes obtained from a canine and the manual tracings obtained from an independent observer are as shown in Fig. 7.7(b). The segmentation results obtained after the incorporation of the learned in-region cost term, on-surface cost term and a combination of the two is as shown in Figs. 7.7(c)-(e), respectively.

7.4 Discussion and Conclusion

Machine-learning based approaches and graph-theoretic approaches [21, 27, 55] proposed for the segmentation of retinal surfaces in SD-OCT images have their own

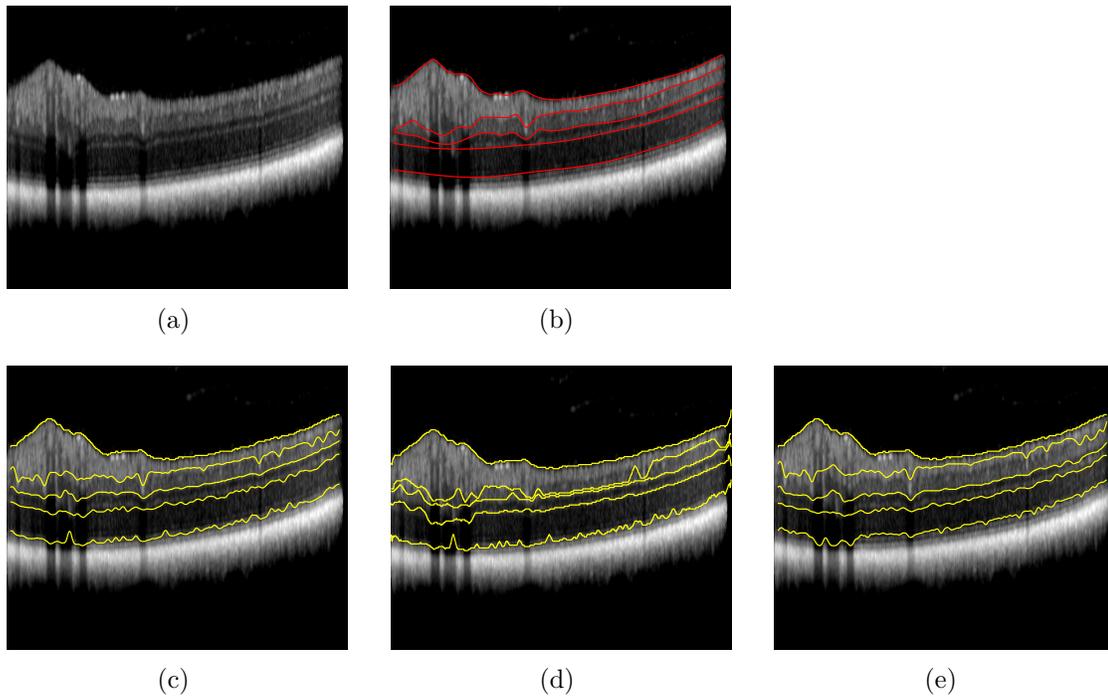


Figure 7.7: Segmentation results obtained on the dog dataset. (a) A slice from an SD-OCT image obtained from a canine alongside (b) the manual tracings obtained from an independent observer and the automated results obtained using (c) the learned in-region cost term, (d) on-surface cost function, and (e) a combination of both.

distinct advantages. Machine-learning based approaches are easily adapted to new models while graph-theoretic approaches are able to segment multiple surfaces simultaneously in 3-D while ensuring the global optimality of the solution with respect to a cost function. Here, we described a framework that combines a machine-learning based technique and a graph-theoretic approach [23] for the automated segmentation of retinal surfaces in SD-OCT images. The machine-learning based automated design of the cost function allowed for the method to be quickly adapted to new models such as animal retinae. The impact of the incorporation of the learned components final segmentation accuracy was also assessed, thereby allowing for the method to be tailored to specific applications as its only requirement is a training set.

When using a machine-learning based approach, the choice of classifier can have a significant impact on the final result. In our previous work (see Chapter 5) we used a k -NN classifier, which performed acceptably well. However, this classifier has known issues associated with large class size disparity and the computation time grows with the dimensionality and size of the training set. Support vector machines (SVM) was one of the options we considered, however the large number of parameters that needed to be optimized proved to be a cumbersome process. Furthermore, the best models that were found were still extremely sensitive to small parameter changes, indicating that the SVMs were likely being over-trained. However, more general models did not provide results comparable to those obtained when using the k -NN classifier. The random forest proved to be a good alternative, that did not have many parameters to optimize and once the random forest was created, classified input vectors very quickly.

As the cost function design relies on the training set, we observed that the difficulty of the manual tracings had an impact on the ability to learn the costs for the regions and surfaces of interest. This difficulty is apparent in the human dataset, where all the images were acquired from subjects that presented with glaucoma. The contrast

in these images was quite low and the difficulty associated with confidently identifying the surfaces is indicated in the high inter-observer variability. While the method was able to differentiate between the regions quite well, learning the surfaces proved to be quite difficult on this set. This is also reflected in the final recommendation of the framework, where the incorporation of the learned in-region cost alone (Experiment I) showed the best overall results.

Similar results were noted in the canine dataset, where the scans were also largely obtained from disease-affected animals. These images were quite anisotropic, with few slices acquired along the y -axis, making the inclusion of 3-D features or contextual information in the method difficult. The images also often showed low contrast making the development of a baseline method that would provide adequate accuracy extremely difficult to design. The low contrast also affected the ability to accurately trace the inner surfaces, and subsequently confounded the learning of the surface properties. Here, again the framework recommended the incorporation of the learned in-region cost term. However, certain surfaces fared well with the combination of the two learned cost terms, indicating that perhaps the learned components could be included for certain surfaces and excluded for other surfaces that are difficult to learn. However, this is quite difficult to gauge, given the numerous interactions that exist when the surfaces are segmented simultaneously. This will, however, be investigated in our future work. It is also important to note that these images were acquired on a Spectralis SD-OCT scanner where the voxel depth is $3.86 \mu\text{m}$. Thus, the overall error is on the order of 2-4 voxels, which is comparable to the results obtained on the other two sets of images.

In the case of the mice scans, the framework recommended the use of a combination of the learned in-region and on-surface cost terms. Although, certain surfaces, such as the outer surfaces, seemed to not need the learned cost terms, the inclusion of the cost terms had a significant impact on the segmentation accuracy of the NF+GCL.

This structure is of importance as it has been shown to correlate with disease [92]. As the NF+GCL thins and becomes difficult to discern, we expect the inclusion of the learned cost terms, especially the in-region cost term, to have a significant impact on the segmentation accuracy.

The graph-theoretic method is a natural choice where numerous regions (with possible similarities) need to be segmented, and the automated machine-learning based cost function design allows for this method to be easily adapted to numerous applications. In particular, the automated machine-learning approach requires less domain (animal/model specific) knowledge on the part of the algorithm development team. Furthermore, as SD-OCT imaging finds increasing use in research studies involving different animal models, we expect that this adaptable method could prove to be an extremely useful tool.

CHAPTER 8

INCORPORATION OF LEARNED SHAPE PRIORS INTO A GRAPH-THEORETIC APPROACH WITH APPLICATION TO THE 3-D SEGMENTATION OF INTRARETINAL SURFACES IN SD-OCT VOLUMES OF MICE

In addition to automatically designing the in-region and on-surface cost terms utilized by the graph-theoretic approach, the cost function can include an additional shape term as described in Section 3.1. As demonstrated in Chapter 7, the inclusion of the learned on-surface cost term did not prove beneficial when segmenting the human glaucomatous data or the dog data. The learned in-region cost term, did however, significantly improve the segmentation accuracy. Thus, before incorporating the shape priors into the graph-theoretic method, its impact on the segmentation result needs to be carefully evaluated.

In this chapter, a graph-theoretic approach that incorporates learned shape priors is presented. The method is applied to the segmentation of 10 retinal surfaces (shown in Fig. 8.1) in SD-OCT images obtained from mice. These densely acquired scans have a higher resolution than routinely acquired clinical scans and are extremely rich in contextual information, making them ideal for learning shape priors. The individual coefficients on the surfaces were optimized as some surfaces may prefer different weighting coefficients. A comparison of this proposed method was conducted to a baseline method [91] (see Chapter 6) that did not incorporate any learned features as well as to the method that incorporated learned in-region and on-surface cost terms and produced the smallest segmentation error (see Chapter 7).

8.1 Method

Recall from Section 3.1, the cost function C_T , when incorporating a shape term can be expressed as:

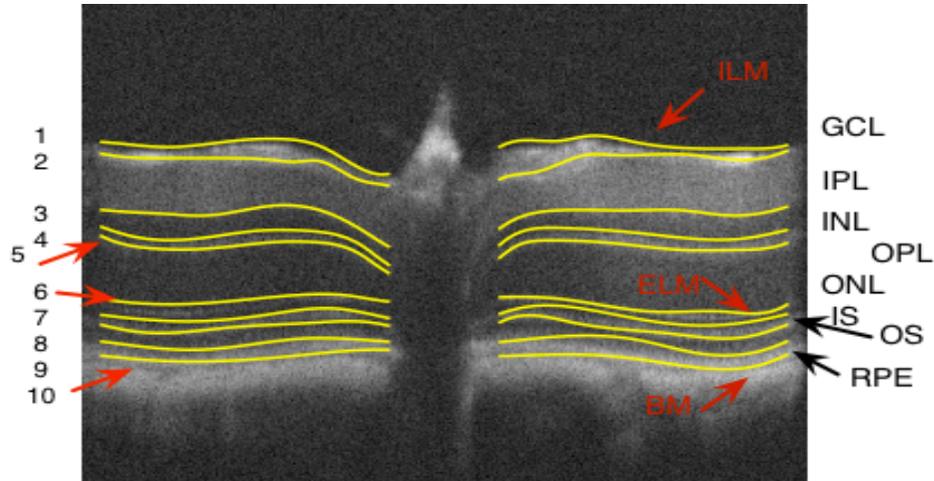


Figure 8.1: A central slice from an SD-OCT volume of a mouse showing the 10 surfaces segmented.

$$C_T = \sum_{i=1}^n (C_I + C_{shape}), \text{ where } C_{shape} = \sum_{i=1}^n \sum_{p,q \in \mathcal{N}} f_{p,q}((\mathcal{S}_i(p) - \mathcal{S}_i(q)) - m_i(p,q)), \quad (8.1)$$

where, C_I represents the cost term derived from image features and $m_i(p, q)$ represents the mean shape prior (indicated by the training set). This term can either be designed using image gradients (as detailed in Chapter 6) or can incorporate learned features using the machine-learning based approach described previously in Chapter 7. If a quadratic function of the form $a(k - m)^2$ (where $k = \mathcal{S}_i(p) - \mathcal{S}_i(q)$) is used as the penalty shape function, then the total shape penalty cost can be represented by:

$$C_{shape} = \sum_{i=1}^n a_i \sum_{p,q \in \mathcal{N}} ((\mathcal{S}_i(p) - \mathcal{S}_i(q)) - m_i(p,q))^2, \quad (8.2)$$

where a_i represents the weighting coefficient for the surface i .

The overall process is outlined in Fig. 8.2. The method begins with the segmen-

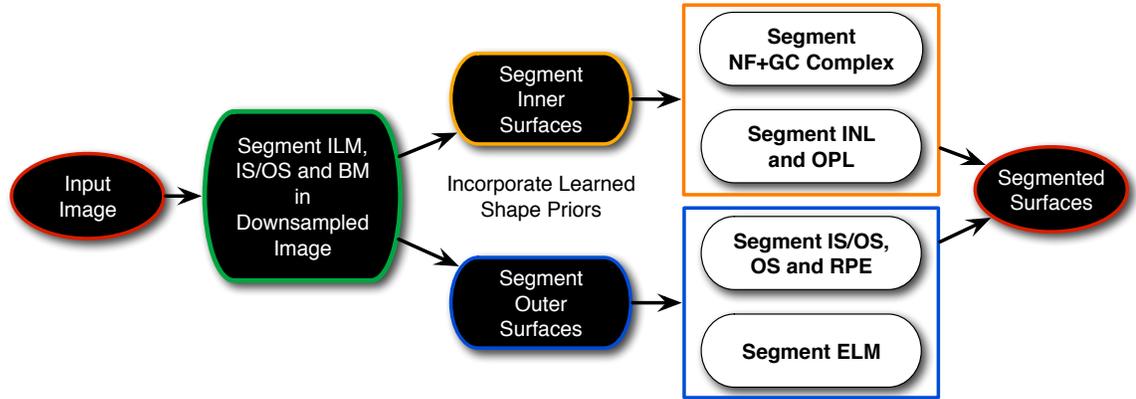


Figure 8.2: Overview of the segmentation process used to segment 10 retinal surfaces in the mice SD-OCT scans. The average location of the inner and outer surfaces are first detected by segmenting the surfaces in a downsampled volume. The inner and outer surfaces are then segmented separately as shown above.

tation of the ILM, IS/OS junction and Bruch's membrane in a downsampled (by a factor of 8) volume in order to obtain an approximate location of the retinal surfaces. This step did not incorporate the learned shape priors and is identical to the first segmentation step of the baseline approach (described in Section 6.1). The location of the IS/OS line is used to create a volume of interest that is then used to segment surfaces 6 - 10 (see Fig. 8.1). All the surface segmentations from here on were done in the full resolution and the learned shape priors were incorporated into the graph-theoretic method. The value of $m_i(p, q)$ was set to the mean change in the topology of surface i as indicated by a training model. The training set utilized and the optimization of the parameter a_i are described further in Section 8.2. The cost function used, as previously described, contained a shape term and an image term C_I , where the latter was either designed using difference of Gaussians filters or incorporated learned texture features (also further described in Section 8.2).

The segmentation of the outer surfaces began with the simultaneous segmentation of the IS/OS line, the bottom of the OS, and the bounding surfaces of the RPE. The ELM was segmented in a second step. The inner five retinal surfaces were segmented

in a similar fashion, beginning with the segmentation of the bounding surfaces of the NF+GC complex. The prior segmentation of the ILM is used to create a small volume of interest within which the ILM and the bottom surface of the NF+GC complex is segmented. The final result of the ILM and the ELM from the previous step are used to create another volume of interest, within which the bounding surfaces of the INL and the OPL (three surfaces) were segmented simultaneously.

8.2 Experimental Methods

The data used to evaluate this method consisted of 10 OCT volumes from 10 normal mice obtained on an SD-OCT Bioptigen (Bioptigen Inc., Durham, NC) scanner (described in Section 7.2.1). In addition to the manual tracings obtained from a retinal specialist, a second bronze standard of tracings were also obtained on the entire volume of all 10 datasets, where the volumes were first segmented using a baseline approach (see Chapter 6) [91] and subsequently, manually corrected. This set was used to learn the shape priors as well as the varying feasibility constraints used in the graph-theoretic method. The shape prior $m_i(p, q)$ was set to the mean of the change noted between columns p and q in the bronze standard for surface i . The varying hard smoothness constraints $\Delta_{p,q}^u$ and $\Delta_{p,q}^l$ were set to $m_i(p, q) \pm 2.6 \times \sigma_i(p, q)$, respectively. A cross-validation scheme was used to learn the shape priors $m_i(p, q)$, where the eight datasets were used as the training set and the remaining two were used as the test set. This process was repeated four times.

As the various surfaces may prefer different values of a_i , the weighting parameter, this value needed to be optimized. A variety of values of $a = \{0.9, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 5.0\}$ were tested and the unsigned border position error for each of the 10 surfaces were noted. Small increments were used to refine the value a_i when an estimate for the surfaces were found.

Once, the value of a_i was optimized, the impact of the shape priors on the segmentation accuracy was gauged. For this, the cost term C_I was designed using the

following two ways:

1. The cost term C_I consisted of image gradients obtained by filtering the image with difference of Gaussians filters as described in baseline method (see Section 6.1). The only learned component in this configuration was the shape priors (Approach I).
2. C_I was designed as a weighted combination of a learned in-region and on-surface cost term as recommended by the automated cost function design framework (described in Section 7.1). This configuration therefore, consists of a learned in-region, on-surface and shape cost term (Approach II).

Both of the above configurations of the graph-theoretic method were compared to the baseline approach (Chapter 6) and the configuration of the method that incorporated the learned in-region and on-surface cost terms (described in Chapter 7).

8.3 Results

8.3.1 Optimization of Weighting Parameter a

The segmentation accuracy was gauged by computing the signed and unsigned border position error (in microns) between the automated segmentation results and the manual tracings. The unsigned border position error noted in the 10 retinal surfaces for the values of a are summarized in Table 8.1. Fig. 8.3 (a) and (b) show the error noted in the inner and outer surfaces, respectively for each of the values of a tested. The value of a that minimized the unsigned border position error for each surface was then used to segment the retinal surfaces to give us the final result.

8.3.2 Segmentation Accuracy

A cross-validation scheme was used when incorporating the learned shape priors where, the expected topology (mean and standard deviation of the surface change between two adjacent locations) of the 10 retinal surfaces were learned from eight

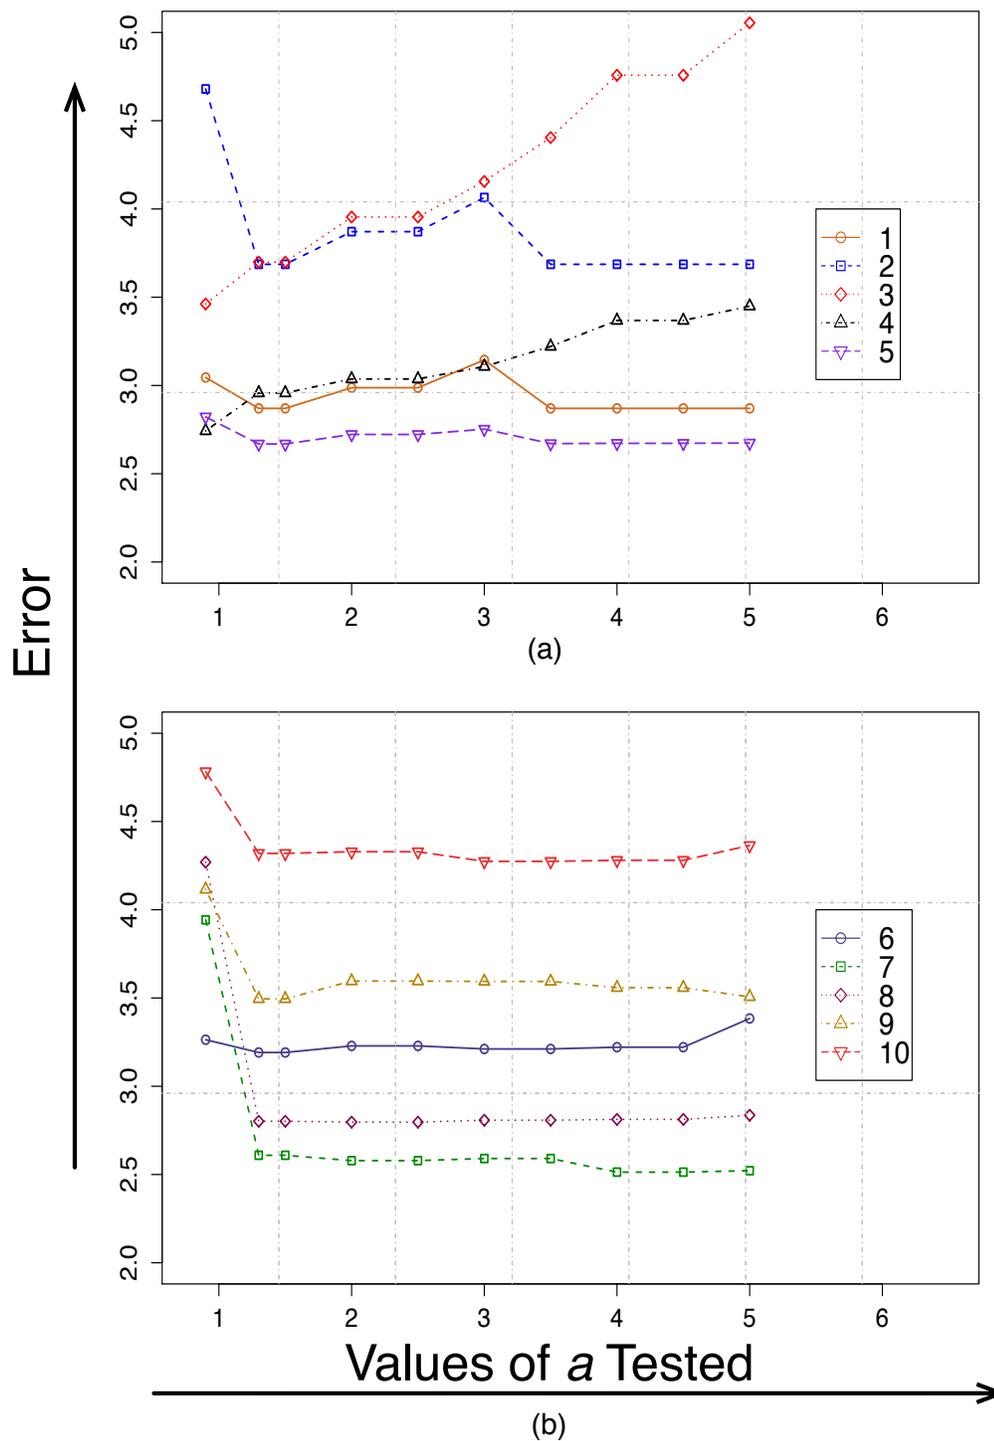


Figure 8.3: Graph of error vs. tested values of a . Plots of the unsigned error against the values of a that were tested for the (a) five inner surfaces and the (b) five outer surfaces.

Table 8.1: Unsigned border position error observed for the 10 retinal surfaces for the range of value of a tested.

Surfs.	0.9	1.3	1.5	2.0	2.5	3.0	3.5	4.0	4.5
1	3.05	2.87	2.87	2.99	2.99	3.15	2.87	2.87	2.87
2	4.68	3.69	3.69	3.87	3.87	4.07	3.69	3.67	3.67
3	3.46	3.70	3.70	3.96	3.96	4.16	4.42	4.76	4.76
4	2.74	2.96	2.96	3.04	3.04	3.11	3.22	3.37	3.37
5	2.82	2.67	2.70	2.72	2.72	2.75	2.67	2.67	2.67
6	3.26	3.19	3.19	3.23	3.23	3.21	3.21	3.22	3.22
7	3.94	2.61	2.61	2.58	2.58	2.59	2.59	2.51	2.51
8	4.27	2.80	2.80	2.80	2.80	2.81	2.81	2.81	2.81
9	4.12	3.50	3.50	3.60	3.60	3.59	3.59	3.56	3.56
10	4.78	4.32	4.32	4.33	4.33	4.27	4.27	4.28	4.28

datasets and subsequently tested on the remaining two datasets. This was repeated four times. The signed and unsigned border position errors for the 10 retinal surfaces is summarized in Table 8.2. The overall signed and unsigned segmentation error for the 10 retinal surfaces was found to be $-0.13 \pm 1.22 \mu\text{m}$ and $3.20 \pm 0.56 \mu\text{m}$, respectively when using Approach I. The errors noted for surfaces 2, 8 and 9 were found to be significantly ($p < 0.005$) smaller in comparison to the results obtained using the baseline method. These results were however, not significantly different from those obtained when using the learned in-region and on-surface cost terms (labeled ‘Learned C_I ’ in Table 8.2). The overall signed and unsigned error obtained when using Approach II was $0.03 \pm 1.43 \mu\text{m}$ and $3.43 \pm 0.58 \mu\text{m}$, respectively. These results were not significantly different from those obtained when using only the learned in-region and on-surface cost terms or Approach I. The errors noted for surfaces 2, 8 and 9 were however, significantly ($p < 0.005$) smaller from those obtained using the baseline approach.

Fig 8.4(a) and (b) show near central slices from two mouse SD-OCT scans. The manual tracings obtained from the retinal specialist overlaid on these images are as shown in Fig. 8.4(c) and (d). The segmentation results obtained when using the

Table 8.2: Accuracy assessment of segmentation results when incorporating learned shape priors. Unsigned and signed border position error[†] (mean \pm SD μm) computed for the two proposed approaches. Previously obtained results from a baseline method (Section 6.3) and the method that incorporated learned texture features into the cost function (Section 7.3.2) are also summarized.

Surfaces	Baseline	Learned C_I *	Approach I $^\circ$	Approach II $^\bullet$
1	2.48 ± 0.72	2.48 ± 0.72	2.91 ± 1.44	2.88 ± 0.80
2	6.97 ± 2.20	4.45 ± 0.56	3.68 ± 0.74	3.94 ± 0.92
3	2.86 ± 0.55	3.24 ± 0.83	3.55 ± 0.74	3.41 ± 0.71
4	2.80 ± 0.73	2.59 ± 0.71	2.77 ± 0.70	2.98 ± 0.69
5	3.16 ± 0.97	3.98 ± 1.33	2.67 ± 0.45	3.84 ± 1.05
6	3.08 ± 1.08	3.01 ± 0.86	3.20 ± 1.00	3.52 ± 1.14
7	2.72 ± 0.98	2.89 ± 0.80	2.51 ± 1.10	2.54 ± 1.11
8	3.39 ± 1.20	3.57 ± 0.98	2.83 ± 1.17	3.23 ± 1.15
9	4.50 ± 0.97	3.58 ± 1.17	3.57 ± 1.13	3.97 ± 1.31
10	3.87 ± 1.36	3.72 ± 1.15	4.29 ± 1.44	4.51 ± 1.76
Mean	3.58 ± 1.33	3.35 ± 0.62	3.20 ± 0.56	3.43 ± 0.58

[†] Errors are computed within an annular region defined by two circles 0.2 and 1.2mm in diameter.

* This approach incorporated the learned texture features into the cost term C_I and was presented earlier in Section 7.3.2.

$^\circ$ This approach incorporated learned shape priors and the C_I is derived from image gradients.

$^\bullet$ In addition to the learned shape priors, C_I was a weighted combination of learned in-region and on-surface cost terms.

baseline approach and after the incorporation of the learned shape priors are as shown in the third and fourth rows, respectively. The learned shape priors were found to be particularly useful near blood vessels and overall, produced a smoother segmentation result.

8.4 Discussion

The graph-theoretic approach formulation described here, incorporated soft shape constraints into the graph which, was observed to produce more accurate and smoother

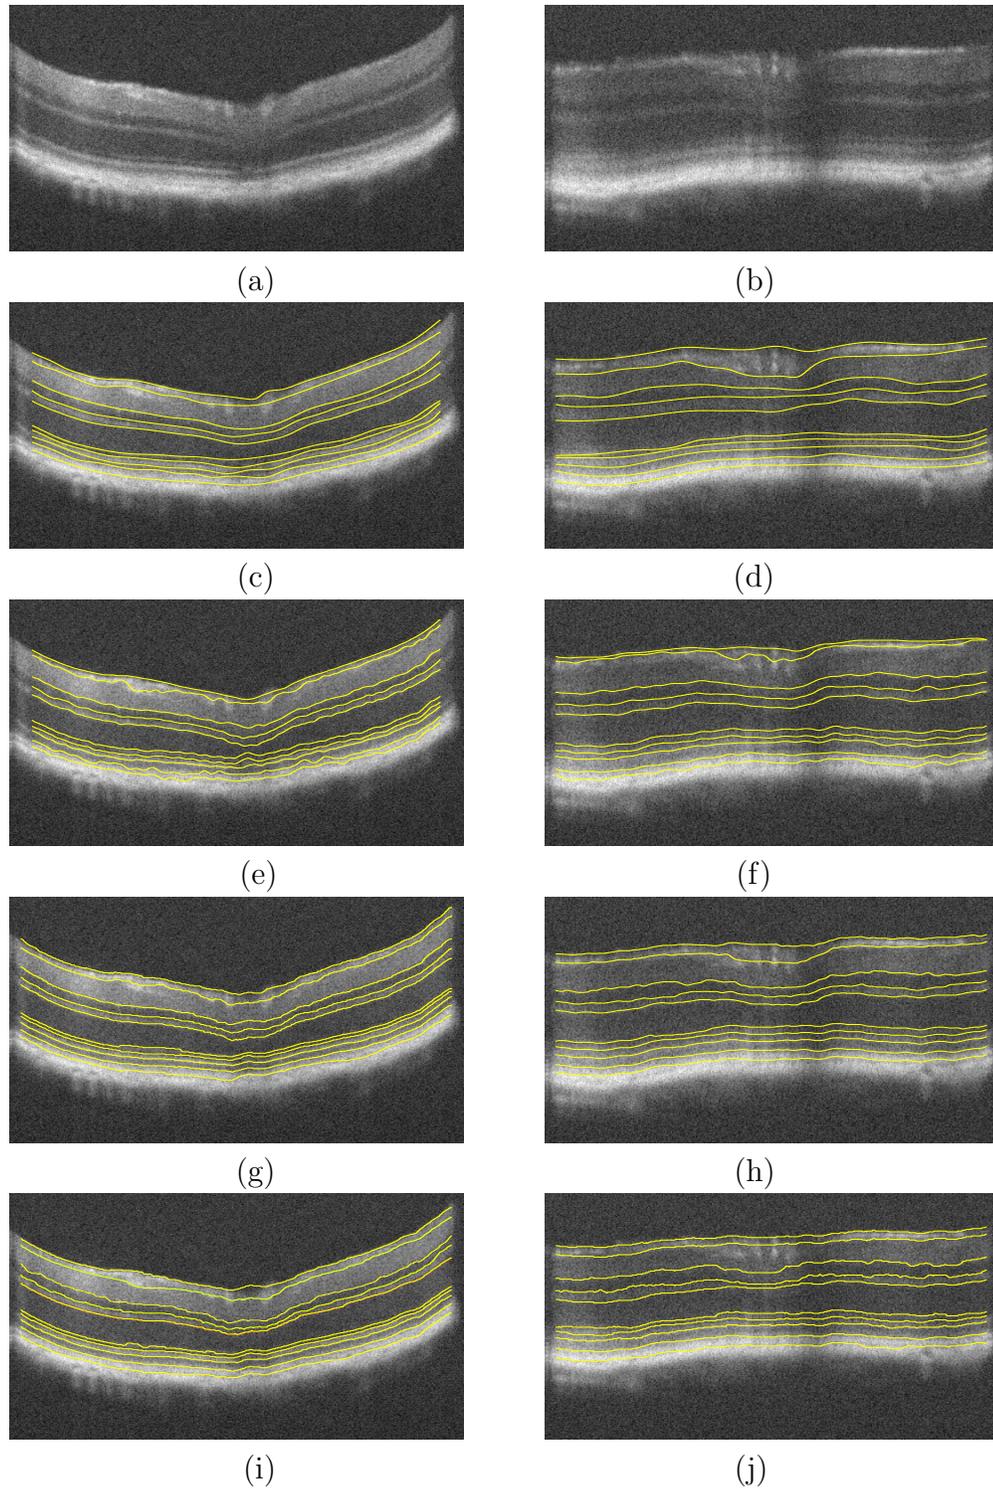


Figure 8.4: Segmentation results obtained after incorporating learned shape priors. (a),(b) Near central slices from two mice datasets with (c),(d) the manual tracings obtained from the retinal specialist overlaid. The segmentation results obtained using (e),(f) the baseline method, (g),(h) Approach I and (i),(j) Approach II.

surfaces. While the incorporation of these learned shape constraints did produce significantly better results with respect to the baseline approach that did not incorporate any learned information, the method did not fare better than when the on-surface and in-region cost terms were designed using a machine-learning approach.

The inclusion of the learned shape constraints is known to substantially increase the number of edges in the graph, which in turn has an impact on the run-time. The mice datasets in particular, are quite large (about 319 MB), containing approximately 4 times the voxels contained in scans typically acquired from humans (about 39 MB) in a clinical setting. While the baseline approach and the approach described in Section 7.1 did segment up to five surfaces simultaneously, the same was not possible here when the learned shape priors were included and the methods had to be suitably altered in order to compensate for the memory limitation. Thus, the incorporation of each new learned component into the graph-theoretic method needs to be weighed against the impact of the component on the segmentation accuracy. While this method did not fare better than when the other terms of the cost function incorporated learned features, this method had a smaller overall running time as the texture feature generation can be an extremely time-consuming step. Furthermore, a smaller graph could be constructed as described by Dufour et al. [57], which would not only decrease the running time further but would also reduce the memory constraints, potentially making this configuration of the graph-theoretic method ideal in situations where speed is a concern.

CHAPTER 9

A GRAPH-BASED APPROACH FOR THE SEGMENTATION OF MULTIPLE SURFACES WITH A SHARED HOLE

As presented in the previous chapters, the graph-theoretic approach can be an extremely useful tool for the simultaneous segmentation of multiple surfaces. However, in some applications, a set of interacting surfaces has a shared “hole” in which the surfaces do not exist. An example of such an application is that of segmenting the neural canal opening (NCO, the boundary of a “hole”) and the surrounding surfaces within spectral-domain optical coherence tomography volumes (Fig. 9.1). The NCO is an important structure relevant to glaucoma [28, 41, 44] as it can provide a stable reference by which structural changes of the optic nerve head [41, 44] can be monitored.

The neural canal opening (NCO) as described previously in Section 2.1, is a three-dimensional structure that marks the anatomic entrance to the neural canal. As this structure shows little to no change in the presence of early glaucoma, it has been proposed that the NCO can be used to as a stable reference plane from which 3-D measurements of the rim and cup can be made [41, 44]. While this structure can be seen in fundus photographs, the border tissue can in some cases, make it difficult to identify the true location of the optic disc boundary. In volumetric SD-OCT images, however, the NCO and the border tissue [93] can be visualized.

Boyer *et al.* [94] described a method that used a parabolic model to segment the optic disc from time-domain OCT (TD-OCT) images. The method also segmented a single surface of the RPE-complex, which was used to find the cup margin. One of the only previously reported approaches for the segmentation of the NCO and surrounding surfaces in SD-OCT images was that of Hu *et al.* [28], who approached the problem by first attempting to segment the surfaces using a graph-based approach [21, 23], while ignoring the fact that the surfaces do not actually exist inside of the neural canal opening. A projection image was then constructed in which the projected location of

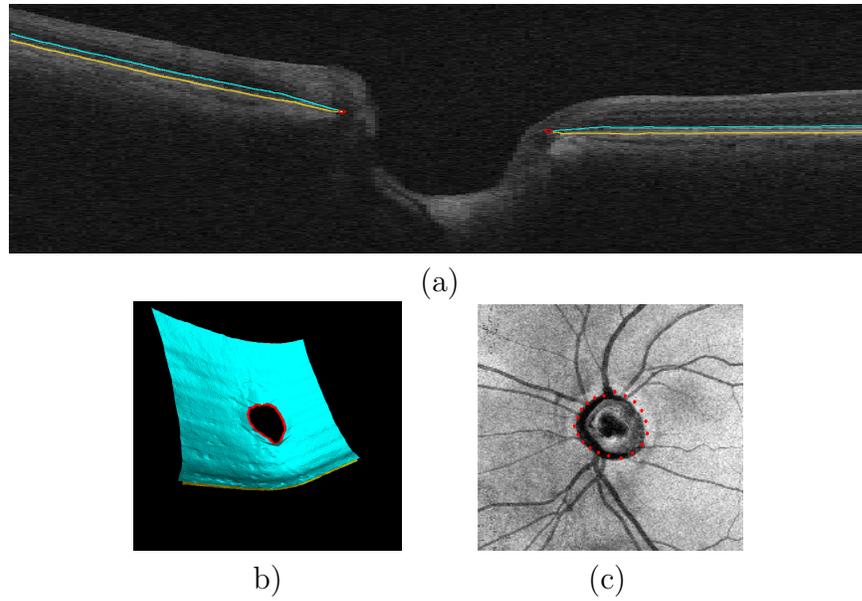


Figure 9.1: Views of the neural canal opening in 3-D and 2-D. (a) A central xz -slice from an SD-OCT scan of a human eye showing the location of the NCO boundary (marked in red). (b) A 3-D rendering of the surfaces and the shared hole (NCO). (c) A 2-D projection image constructed from the SD-OCT image showing the xy -location of the NCO (marked in red).

the neural canal opening (i.e., projected boundary of the hole) is visible. However, the fact that the initial multiple-surface segmentation ignored the existence of the hole makes the segmentation more prone to errors in this region. For instance, a deeper structure inside of the NCO could “pull” the surfaces away from their actual locations near the boundary of the hole (as shown in Fig. 9.2). Furthermore, they were only able to quantitatively evaluate their algorithm in the projected $x-y$ plane using projected manual tracings of the optic disc traced on fundus photographs, and thus, it is difficult to assess the accuracy of the neural canal opening points in the z -direction (i.e., depth direction).

Here, we propose the use of an iterative approach that simultaneously segments the surfaces in the RPE-complex as well as the NCO. In particular, the general framework of our approach reflects an extension of the graph-theoretic approach presented by Li

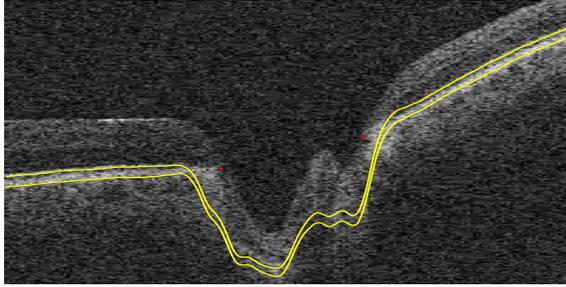


Figure 9.2: Example of surface segmentation errors frequently seen at the NCO. Here, the surfaces are pulled away from the true locations near the boundary of the NCO towards deeper underlying structures.

et al. [23] for simultaneously segmenting multiple surfaces (without a hole) to be able to segment multiple surfaces with a hole. We propose extending the cost function to incorporate costs associated with this additional structure (the hole boundary), as well as incorporating new constraints in order to represent the feasibility of the set of surfaces with the shared hole boundary.

9.1 Method

Assume we have a volume of dimensions $X \times Y \times Z$ and wish to find n layered surfaces with a shared hole as illustrated in Fig. 9.1. The presence of the hole, divides the volume into three regions as depicted in Fig. 9.3, namely the outside \mathcal{O} , the boundary \mathcal{B} and the inside of the hole \mathcal{H} . Intuitively, we will consider a surface set with shared hole feasible if individual surface smoothness constraints are satisfied outside the hole, if the surfaces obey minimum and maximum surface distance constraints outside the hole, if the surfaces come together at the hole boundary, and if the hole boundary satisfies its own set of smoothness constraints, as described further in Section 9.3.1. Given appropriate cost functions (as defined in Section 9.3), we determine a feasible set of surfaces with a shared hole using an iterative approach (see Algorithm 1).

This iterative approach as applied to the segmentation of two surfaces and the

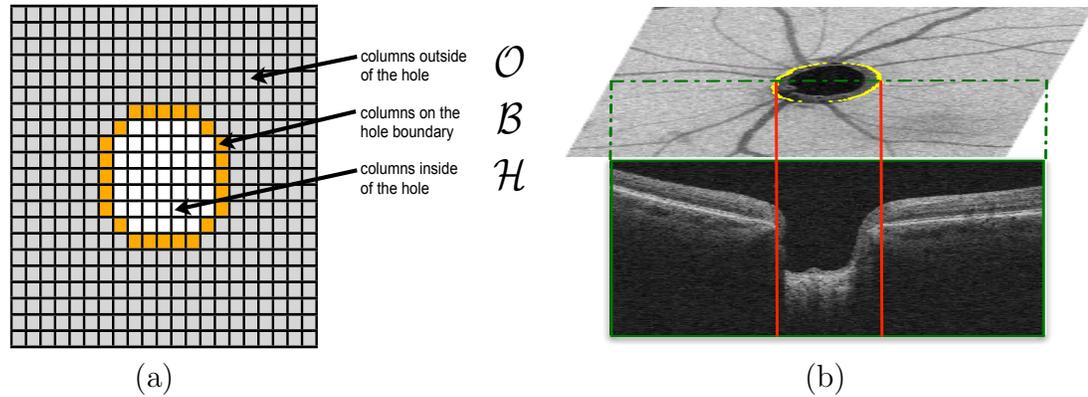


Figure 9.3: Representation of three regions created by the ‘surfaces + hole’ configuration. (a) The 2-D overview of the division of columns into the three sets: the hole boundary, the hole itself and the region outside the hole. (b) The 2-D segmentation of the NCO boundary will be used to define set \mathcal{B} .

Algorithm 1 Segmentation of Multiple Surfaces with Shared Hole

input: initial segmentation of set of surfaces, cost function
 create projection image using initial surface segmentation
 compute initial estimate of projected hole boundary columns in the projection image (Section 9.2)
while not converged **do**
 compute optimal set of feasible surfaces with given hole estimate (Section 9.3)
 update projection image
 compute projected hole boundary columns in the projection image
end while

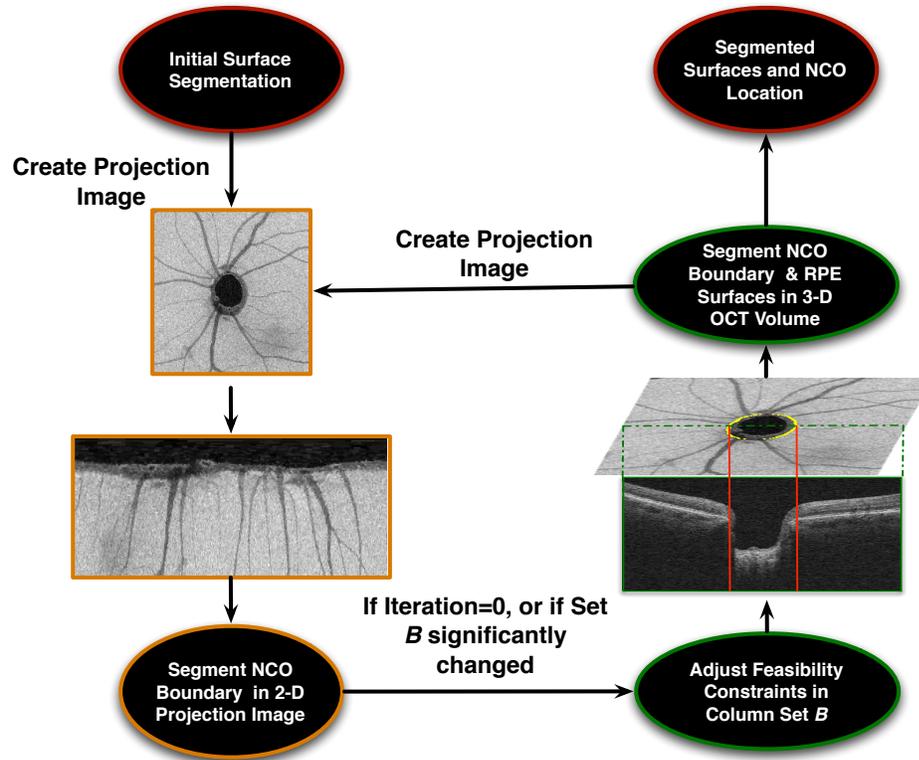


Figure 9.4: Schematic showing the segmentation process that includes a 2-D and 3-D component.

NCO is also illustrated in Fig. 9.4. The method begins by creating a projection image in which the projected boundary columns of the NCO are segmented (see Section 9.2) by finding the minimum-closure in a graph. Next, given this “fixed” set of hole boundary columns \mathcal{B} , we find the corresponding optimal set of feasible surfaces that meet at the hole boundary in the volumetric image by solving another single minimum-closure problem in a constructed graph (Section 9.3). The segmented surfaces are then used to update the projection image, followed by the update of the boundary column-set \mathcal{B} that is obtained by segmenting the projection image. We repeat this process until we achieve convergence of the segmented boundary column-set \mathcal{B} .

9.2 Segmenting the Projected Hole Boundary Columns

We begin with an initial segmentation [25] of three surfaces, namely, the inner limiting membrane (ILM), the junction of the inner and outer segments (IS/OS line) of the photoreceptors and the Bruch's membrane (BM) in the volumetric image. In order to find a smooth representation of a surface that cuts through the optic nerve disruption, a thin-plate spline (TPS) [71] is fit to the segmentation of the Bruch's membrane. The projection image P , is then created using this interpolated smooth surface, such that

$$P(x, y) = \max \{I(x, y, z) | z \in (\mathcal{S}(x, y) - 20, \mathcal{S}(x, y) + 20)\} \quad (9.1)$$

where $\mathcal{S}(x, y)$ is the surface obtained after fitting the TPS to the IS/OS line.

This projection image was then polar transformed (see Figs. 9.5(a) and (b)) so that the NCO boundary appears as a line in the image instead of a circle, and was segmented using a graph-theoretic approach similar to our previous work [28]. The formulation of the graph-theoretic approach used here, however, incorporated shape priors [30] that further influenced the surface topology and provided a smooth surface. The cost function design did not incorporate vessel information, but utilized a stationary wavelet decomposition (SWD) step [95] to filter the strong gradients seen at the vessel shadows. The filtered image (see Fig. 9.5(c)) was obtained by suppressing the horizontal and diagonal details when reconstructing the image. A gradient map was obtained by passing the filtered image through a difference of Gaussians [96] filter, which was then used as the cost function.

9.3 Simultaneous Segmentation of the Surfaces and Their Shared Hole

We will consider a surface set with a shared hole feasible if:

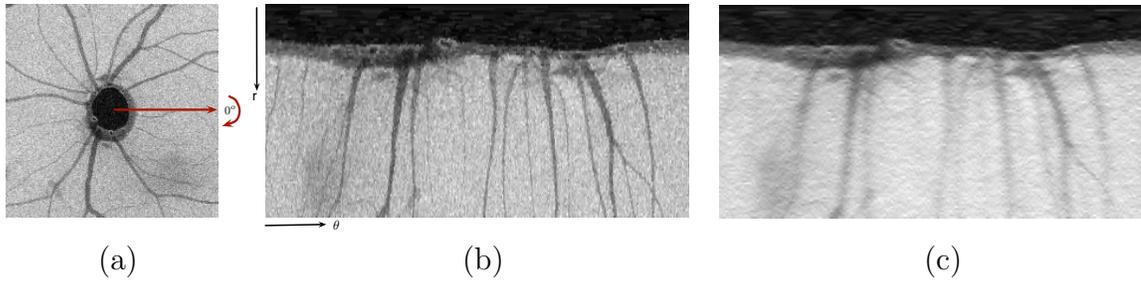


Figure 9.5: Result of stationary wavelet transform-based filtering. (a) The projection image is created from a small region around the IS/OS line and then transformed into the (b) polar coordinate space. This image is then decomposed using a Haar SWD and the response from the vertical and diagonal details are suppressed to give us the final (c) filtered result.

1. individual surface smoothness constraints are satisfied outside the hole,
2. the surfaces obey minimum and maximum surface distance constraints outside the hole,
3. the surfaces come together at the hole boundary, and
4. the hole boundary satisfies its own set of smoothness constraints, as described later in this section.

The formal definition of these constraints and the cost function utilized by the graph-theoretic approach are described in further detail below.

9.3.1 Feasibility Constraints

As shown in Fig. 9.3, we assume that each (x, y) -column can be divided into one of three categories:

1. columns for which each of the n surfaces will intersect the column at exactly one z -location (i.e., columns outside the hole \mathcal{O}),
2. columns for which the boundary of the hole exists at exactly one z -location (i.e., columns on the boundary of the hole \mathcal{B}), and

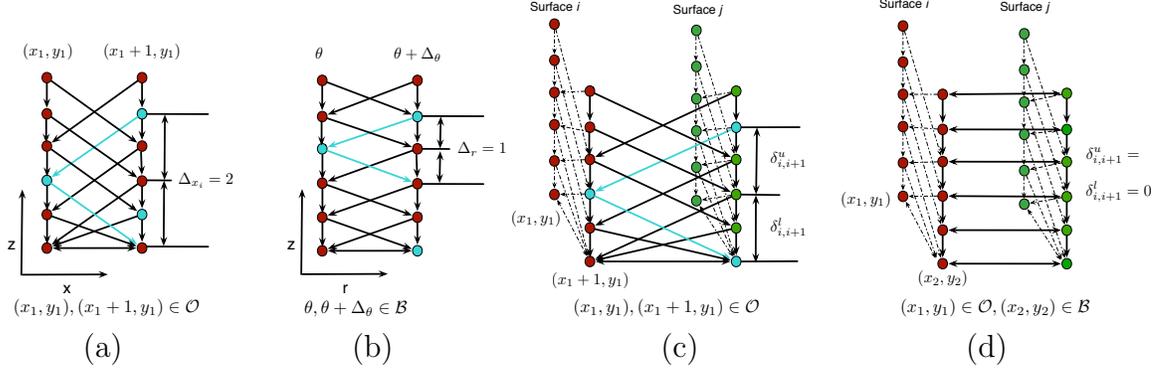


Figure 9.6: Constraints used for regions \mathcal{B} and \mathcal{O} . The illustration show the smoothness constraints within regions (a) \mathcal{O} and (b) \mathcal{B} , and the surface-interaction constraints within regions (c) \mathcal{O} and (d) \mathcal{B} .

3. columns for which the surfaces are not defined because of the hole (i.e., columns inside the hole \mathcal{H}).

The n surfaces $\{f_1(x, y), \dots, f_n(x, y)\}$ are defined over the columns in $\mathcal{O} \cup \mathcal{B}$, where the smoothness constraint for each surface i is defined by:

$$|f_i(x+1, y) - f_i(x, y)| \leq \Delta_{x_i} \quad \text{and} \quad (9.2)$$

$$|f_i(x, y+1) - f_i(x, y)| \leq \Delta_{y_i}, \quad (9.3)$$

where Δ_{x_i} (Δ_{y_i}) is the x -smoothness (y -smoothness) parameter for surface i (see Figs. 9.6(a) and (b)). For each of the columns in \mathcal{O} , each i - $(i+1)$ pair of surfaces (where surface $i+1$ is directly “above” surface i) is also constrained by the surface distance constraints defined as:

$$\delta_{i,i+1}^l \leq f_{i+1}(x, y) - f_i(x, y) \leq \delta_{i,i+1}^u, \quad (9.4)$$

where $\delta_{i,i+1}^l$ and $\delta_{i,i+1}^u$ is the minimum and maximum allowed distance between the surfaces, respectively. Note that these two feasibility constraints are defined similarly

as in the standard multiple surface segmentation problem (without a hole) [23].

In addition, we require that the x - y projection of the hole boundary be representable using a function defined in polar coordinates $f_b(\theta)$ (i.e., having one intersection per sampled angular ray) and that this projected boundary be sufficiently smooth:

$$|f_b(\theta) - f_b(\theta + \Delta_\theta)| \leq \Delta_r , \quad (9.5)$$

where Δ_θ is the angular distance between sampled rays and Δ_r is the smoothness parameters specifying the maximum change in radial position between angular rays. Furthermore, for columns on the shared boundary of the hole (\mathcal{B}), we require that the minimum and maximum distance between all surfaces be equal to 0 (i.e., the surfaces come together at any column on the shared hole boundary), as depicted in Figs. 9.6(c) and (d).

The constraints for regions $\mathcal{B} \cup \mathcal{O}$ can also be expressed in polar coordinates where the volumetric scan be unwrapped along a central axis. Thus, a surface can now be expressed as a function of the radial distance r and the angle θ from the central axis $\mathcal{S}_i(r, \theta) = z$. Within this coordinate system the smoothness constraints of the surfaces within \mathcal{O} can be expressed as

$$|f_i(r + 1, \theta) - f_i(r, \theta)| \leq \Delta_{r_i} \quad \text{and} \quad (9.6)$$

$$|f_i(r, \theta + 1) - f_i(r, \theta)| \leq \Delta_{\theta_i} , \quad (9.7)$$

where Δ_{r_i} and Δ_{θ_i} represent the maximum allowed change in the z direction along the radial and angular axes for surfaces i . The surface interaction constraints do not change, as the surface interaction is not expected to change with a conversion to the polar coordinate system. Thus, the interaction constraints within \mathcal{O} can be expressed as:

$$\delta_{i,i+1}^l \leq f_{i+1}(r, \theta) - f_i(r, \theta) \leq \delta_{i,i+1}^u . \quad (9.8)$$

9.3.2 Cost Function

For the segmentation of the surfaces and the shared hole, every voxel in the volume will be assigned $n+1$ costs: one cost associated with the unlikeliness of belonging to each of the n surfaces and one cost associated with belonging to the boundary of the shared hole. Let $c_{\text{surf}_i}(x, y, z)$ reflect the cost associated with surface i ($i = 1, \dots, n$) and $c_b(x, y, z)$ reflect the cost associated with the hole boundary. Then, the cost of a feasible surface set with a shared hole (given \mathcal{O} , \mathcal{B}) is defined as follows:

$$C_T = \sum_{i=1}^n \sum_{(x,y) \in \mathcal{O}} c_{\text{surf}_i}(x, y, f_i(x, y)) + \lambda \sum_{(x,y) \in \mathcal{B}} c_b(x, y, f_1(x, y)) \quad (9.9)$$

where the first term is associated with the “on-surface” costs in region \mathcal{O} and the second term is associated with the cost of the shared hole boundary points. The assumption that the surfaces are not defined within \mathcal{H} is dependent on an accurate estimation of the boundary column \mathcal{B} . However, should the estimate of \mathcal{B} not be completely accurate, the lack of a cost term associated with the region \mathcal{H} could lead to the surfaces not being segmented correctly. The presence of underlying structures that could dominate the surface segmentation could however, cause errors within region \mathcal{H} to propagate to \mathcal{O} or \mathcal{B} . A suitably weighted on-surface cost term however, could retain information regarding the surfaces within \mathcal{H} while preventing this region from dominating the final solution. Thus, in these cases it may be useful to express the cost function as:

$$C_T = \sum_{i=1}^n \sum_{(x,y) \in \mathcal{O}} c_{\text{surf}_i}(x, y, f_i(x, y)) + \lambda \sum_{(x,y) \in \mathcal{B}} c_b(x, y, f_1(x, y)) + \alpha \sum_{i=1}^n \sum_{(x,y) \in \mathcal{O}} c_{\text{surf}_i}(x, y, f_i(x, y)) \quad (9.10)$$

where the co-efficient α can be set to a value between 0 and 1. We used this formu-

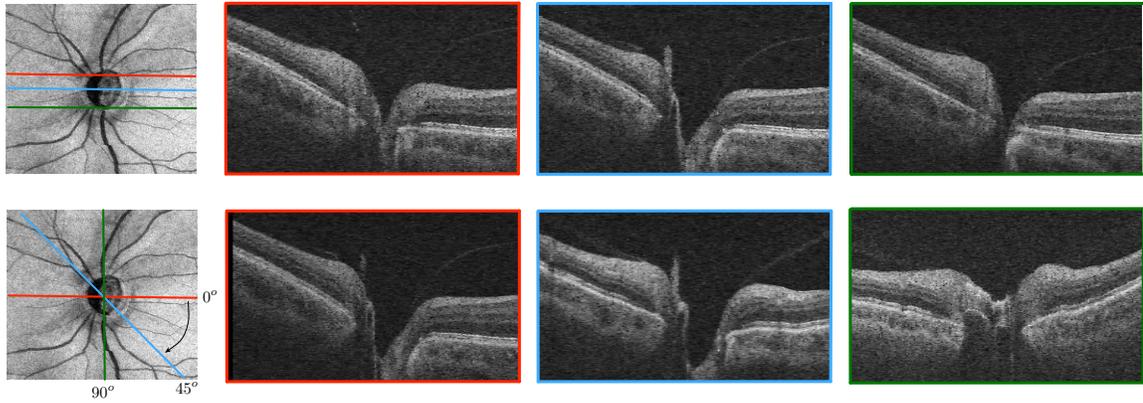


Figure 9.7: Polar transformation of volumetric images. The red, blue and green outlines slices in the top row show B-scan slices that correspond to the locations marked in red, blue and green, respectively on the projection image. The NCO location is difficult to identify in these B-scans, especially close to the upper and lower part of the NCO. The red, blue and green slices in the second row show radial slices at 0° , 45° and 90° , respectively, as marked on the projection image.

lation of the cost function in this work.

Before computing the cost functions for the regions \mathcal{B} , \mathcal{O} and \mathcal{H} , the volumetric image was also polar-transformed with an angular spacing of 1° . In the volumetric images, the NCO can be difficult to identify at the upper and lower regions (marked in red and green in the top row of Fig. 9.7) of the NCO as the end points get closer together. The large blood vessels that characterize these regions of the retina also create large vessels shadows which can make the manual delineation the NCO nearly impossible. However, in the polar-transformed scans, the NCO's continuity is maintained, making it easier to visualize and delineate, as shown in the second row of Fig. 9.7. Another advantage of using the polar-transformed scan is that the boundary column estimate \mathcal{B} computed in the previous step does not need to be converted into Cartesian coordinates and all of the processing can be done in polar coordinates.

The cost function for the regions \mathcal{O} and \mathcal{H} consisted of on-surface cost terms derived from difference of Gaussians [96] filters, while the cost function for \mathcal{B} incorporated textural features learned from a training set. As the NCO boundary can be

modeled as a “corner”, the textural features used to learn the properties of the NCO boundary points included corner detectors such as Harris [97] and SUSAN [98] as well as first order steerable Gaussian derivatives [99] at scales $\sigma = \{1, 2, 3, 4, 5\}$ and orientations $\theta = \{0^\circ, 30^\circ, 60^\circ, 90^\circ, -60^\circ, -30^\circ, -90^\circ\}$. Within the volumetric image $\mathcal{I}(x, y, z)$, recall that the z -axis refer to the depth within the volumetric image. If we further assume that the x -axis refers to the horizontal axis along each slice within the volume and the y -axis refer to the slices, then the the first order 2-D steerable Gaussian filters $G'_{2,\theta}$ is computed as:

$$G'_{2,\theta} = \cos(\theta)G'_{2,0^\circ} + \sin(\theta)G'_{2,90^\circ}, \quad \text{where} \quad (9.11)$$

$$G'_{2,0^\circ} = \frac{\partial}{\partial x}(e^{-(x^2+z^2)}) \quad \text{and} \quad G'_{2,90^\circ} = \frac{\partial}{\partial z}(e^{-(x^2+z^2)}) \quad (9.12)$$

In 3-D, we can use the y -axis slices to provide further context while only differentiating along the x - and z -axes, such that:

$$G'_{3,0^\circ} = \frac{\partial}{\partial x}(e^{-(x^2+y^2+z^2)}) \quad \text{and} \quad G'_{3,90^\circ} = \frac{\partial}{\partial z}(e^{-(x^2+y^2+z^2)}). \quad (9.13)$$

The training set (described further in section 9.4) utilized here consisted of 25 radial volumes that contained 180 slices, where the NCO was identified in all of the slices. Thus, the training set contained 9000 NCO samples. The entire background was not utilized due to the large class size disparity, and a random selection of five background samples from near the NCO was used. A random forest, an ensemble classifier [56], was then trained to identify the NCO boundary points in the volumetric image. The individual feature significances obtained when training the random forest indicated that the steerable Gaussian derivatives were better descriptors of the NCO boundary than the corner detectors. As described previously in Section 7.1.1, features with an individual feature significance less than 2% were excluded from the final feature list, yielding a set of 20 features. The probability map produced by the

classification process was used as the cost term at all column locations in \mathcal{B} .

Two surfaces, the IS/OS line and Bruch's membrane, and their shared hole were then segmented simultaneously in the polar coordinate space using the above defined feasibility constraints and cost functions.

9.4 Experimental Methods

The data used in this experiment consisted of 44 optic nerve head SD-OCT scans obtained from 44 human patients that presented with varying stages of glaucoma. The scans were obtained on a Cirrus (Carl Zeiss Meditec, Inc., Dublin, CA) SD-OCT scanner and were acquired from a region 2mm x 2mm x 6mm and contained 200 x 200 x 1024 voxels. If the subject had been imaged more than once, the first scan acquired was used in this experiment. If the first scan showed substantial acquisition associated artifact such as lateral motion artifact or a large truncation error, then the next available scan was chosen. Fig. 9.8(a) shows the projection image of a scan that was replaced with a repeat scan (see Fig. 9.8b). Such artifacts were observed in 3 out of the 44 subjects.

The volumetric scans were converted into the polar coordinate space where the slices were 1° apart. Manual delineations were obtained from an independent expert (trained to detect the NCO boundary in SD-OCT images) on 10 randomly selected radial slices from each of the 44 scans, giving us a total of 880 traced points in 440 slices. These tracings were then verified by a second independent expert, with a third and final verification being performed by a glaucoma specialist to give us our consensus manual tracings.

The 25 datasets used to train the NCO classifier were also obtained from human patients on a Cirrus SD-OCT scanner, using the same imaging protocol described above. These scans were a subset of the dataset utilized previously [28] for the evaluation of the NCO segmentation. The scans were transformed into the polar coordinate space and the NCO was manually delineated by an independent observer

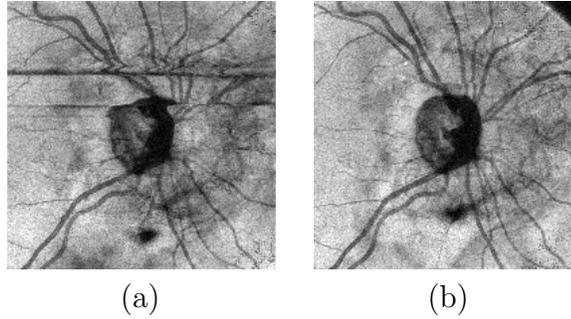


Figure 9.8: Example of scan with substantial artifacts. Projection images from (a) a scan that showed large motion artifacts and (b) the scan that it was replaced with.

in all 180 radial slices.

Previously [28,100], the NCO was segmented using a two-step approach where the NCO boundary located in a projection image was projected down to the segmented Bruch's membrane. In order to gauge the influence of segmenting the surface set and the shared hole simultaneously, the segmentation accuracy obtained using the proposed method was also compared to results obtained when using a two-step approach. In particular, the boundary column-set \mathcal{B} was segmented in the projection images in two ways, namely:

- a pixel classification approach [25] (Approach I), and
- the initial estimate of column-set \mathcal{B} (Approach II) (described in Section 9.2), which is comparable to the method proposed previously [28].

These two boundary column-set estimates were then projected down onto the initial segmentation of the Bruch's membrane to give us the 3-D location of the NCO. As the proposed method is iterative, the errors associated with the first iteration (Approach III - Iter. 1) and the last iteration (Approach III - Final Iter.) were also computed.

The metrics used to gauge the accuracy of the segmentation consisted of the unsigned difference between:

1. the 2-D segmentation in the projection image and the manual delineations,
2. the z locations of the automated segmentation and the manual delineations,
and
3. the 3-D Euclidean distance between the automated segmentation and the manual delineations.

In addition to estimating the segmentation accuracy, it is also important to gauge the ability of the proposed method to measure clinically relevant parameters. The minimum rim width (MRW) has been proposed as a parameter that could be used to gauge the loss of axonal tissue associated with the progression of glaucoma [93]. This parameter, defined as the minimum distance from the NCO to the ILM, was computed using the consensus manual tracings obtained from the experts as well as the automated results obtained using the above mentioned approaches. Statistical analyses (paired t -test comparisons) were conducted in order to gauge if the MRW measurements were significantly different from those obtained from the manual tracings.

9.5 Results

Table 9.1 shows the summary of the accuracy assessment conducted using the three metrics described above. The mean unsigned 2-D error for Approaches I, II and III (final iteration) was found to be $128.36 \pm 61.23 \mu\text{m}$, $62.28 \pm 26.74 \mu\text{m}$ and $46.05 \pm 28.40 \mu\text{m}$, respectively. The 2-D error for Approach II and the first iteration of Approach III do not differ as the same 2-D segmentation result were utilized. The z -depth error for Approaches I, II and the first and last iteration of Approach III were found to be $46.85 \pm 30.21 \mu\text{m}$, $117.21 \pm 43.14 \mu\text{m}$, $28.69 \pm 15.10 \mu\text{m}$ and $29.58 \pm 20.24 \mu\text{m}$, respectively. The 3-D Euclidean distance (in microns) between the manual delineations and the segmented NCO for Approaches I, II and the first

and last iteration of Approach III were found to be $139.67 \pm 61.68 \mu\text{m}$, $136.77 \pm 38.24 \mu\text{m}$, $69.24 \pm 24.12 \mu\text{m}$ and $55.29 \pm 33.97 \mu\text{m}$, respectively.

The errors were also computed in voxels and are presented in parentheses in Table 9.1. The errors noted using the four approaches were also compared statistically. The 95% confidence intervals for the 2-D error observed for Approaches I, II and the final iteration of Approach III were $[110.27 \mu\text{m}, 146.45 \mu\text{m}]$, $[54.38\mu\text{m}, 70.18\mu\text{m}]$ and $[37.66\mu\text{m}, 54.44\mu\text{m}]$, respectively. The errors noted in Approaches I and II were significantly different ($p < 0.01$) from those obtained after the final iteration of Approach III. The 95% confidence intervals for the depth error seen in the four approaches were $[37.86\mu\text{m}, 55.69\mu\text{m}]$, $[104.29 \mu\text{m}, 129.75 \mu\text{m}]$, $[24.18 \mu\text{m}, 33.10\mu\text{m}]$ and $[23.57 \mu\text{m}, 35.51\mu\text{m}]$, respectively. Here, the confidence intervals for the depth error in the first and last iterations of the proposed method overlapped, but the interval found for the final iteration did not overlap with those computed for Approaches I and II. The 95% confidence intervals for the 3-D error for Approaches I, II and the first and last iterations of Approach III were $[121.45 \mu\text{m}, 157.90 \mu\text{m}]$, $[125.47 \mu\text{m}, 148.07 \mu\text{m}]$, $[60.64 \mu\text{m}, 77.85 \mu\text{m}]$ and $[45.25 \mu\text{m}, 65.33 \mu\text{m}]$, respectively. The errors noted in Approaches I and II were found to be significantly larger ($p < 0.001$) than those obtained after the final iteration of Approach III. The errors noted after the first iteration of Approach III was significantly larger ($p < 0.01$) as well.

Fig. 9.9 show the iterative results obtained on a dataset. Figs. 9.9(a) - (c) shows the 2-D results obtained in three iterations (marked in blue, green and yellow). The change in the 3-D location of the segmented NCO is as depicted in Fig. 9.9(d). This slice comes from the radial location marked by the orange arrow. As shown, the error in the location of the NCO substantially reduces from iteration to iteration, until the final result coincides with the manually delineated point.

Figs. 9.10(a) - (b) show the two iterations obtained on another dataset. Here, the initial 2-D segmentation showed an areas with a substantial error as indicated by the

Table 9.1: Summary of accuracy assessment of the NCO segmentation. The errors are expressed in microns and voxels [†].

	2-D	z	3-D
Approach I [*]	128.36 ± 61.23 (4.28 ± 2.04)	46.85 ± 30.21 (23.99 ± 15.47)	139.67 ± 61.68
Approach II [°]	62.28 ± 26.74 (2.08 ± 0.89)	117.21 ± 43.14 (60.01 ± 22.09)	136.77 ± 38.24
Approach III - Iter. 1 [•]	62.28 ± 26.74 (2.08 ± 0.89)	28.69 ± 15.10 (14.69 ± 7.73)	69.24 ± 24.12
Approach III - Final Iter. [∞]	46.05 ± 28.40 (1.54 ± 0.95)	29.58 ± 20.24 (15.15 ± 10.36)	55.29 ± 33.97

[†] Errors computed in voxels is presented in parentheses.

^{*} The 2-D pixel classification approach boundary was projected onto the segmented Bruch's membrane.

[°] A graph-based 2-D segmentation of the NCO was projected onto the segmented Bruch's membrane.

[•] The first iteration of the proposed 3-D method.

[∞] The last iteration of the proposed 3-D method.

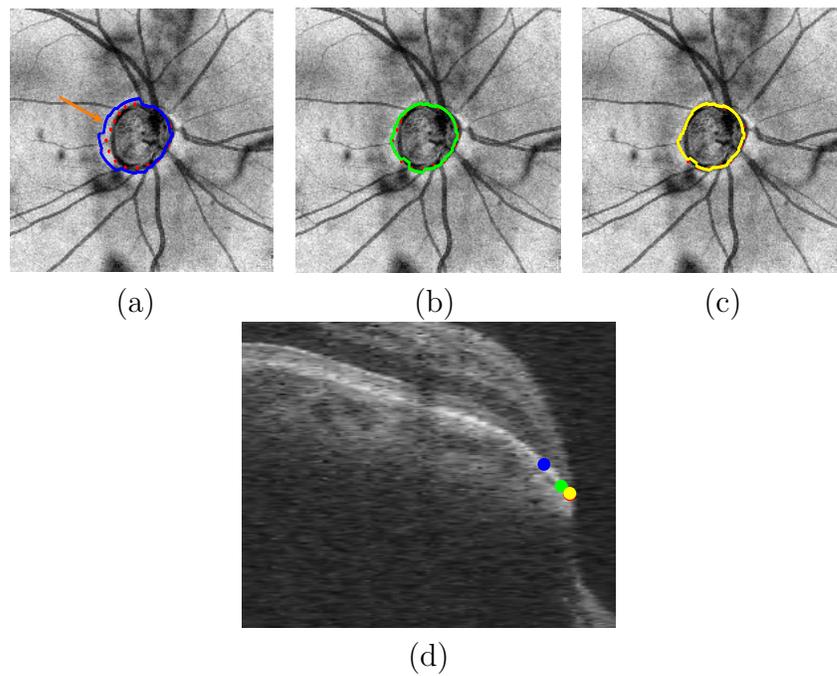


Figure 9.9: A segmentation result. An example of the segmentation result obtained after the (a) first (blue), (b) second (green) and (c) third (yellow) iteration. The manual tracings are marked in red. (d) The 3-D location of the segmented NCO and the manual tracings for the location indicated on (a).

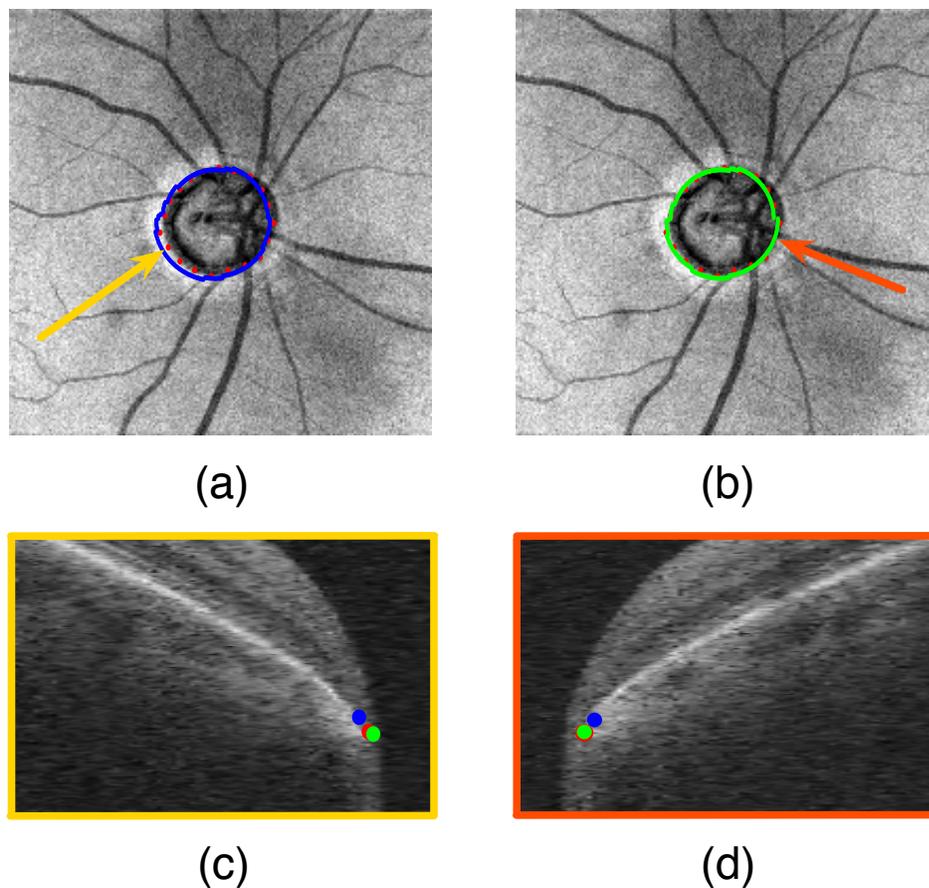


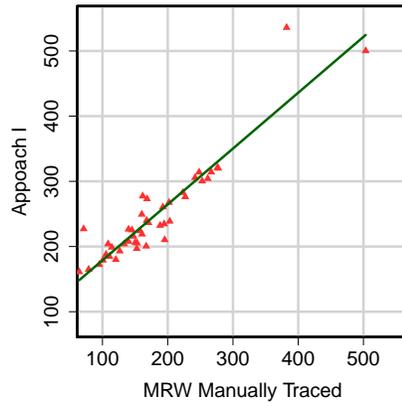
Figure 9.10: Another example of the NCO segmentation. The results obtained from the (a) first (blue) and (b) second (green) iteration. The manual tracings are marked in red. (c), (d) The 3-D location of the segmented NCO and the manual tracings for the locations indicated on (a).

yellow arrow. This error is however, eliminated in the second iteration, as shown in Fig. 9.10(c). In the absence of border tissue, the iterative method converges to the NCO by the second iteration as shown in Fig. 9.10(d). Most datasets converged after two iterations with less than 10% of the images requiring up to five iterations.

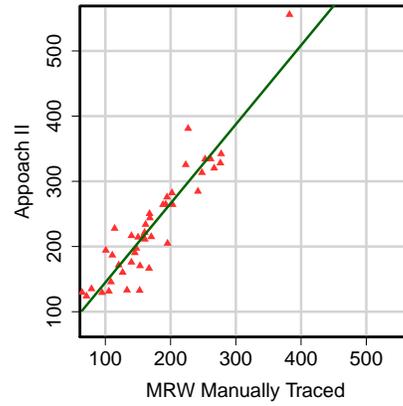
The MRW computed using the manually delineated points was found to be $179 \pm 81.46 \mu\text{m}$. The MRW computed using Approaches I, II and the first and last iterations of Approach III were $247.23 \pm 74.95 \mu\text{m}$, $241.50 \pm 103.94 \mu\text{m}$, $193.37 \pm 77.05 \mu\text{m}$ and $184.09 \pm 84.06 \mu\text{m}$, respectively. The MRW measurements obtained using

Approaches I, II and the first iteration of Approach III were found to significantly ($p < 0.0001$) differ from those obtained using the manual tracings. The MRW computed using the final NCO location obtained using Approach III however, was not significantly ($p = 0.064$) different from those obtained using the manual tracings. The MRW measurements obtained from the four approaches were also correlated to the MRW measurements obtained using the manually delineated points. Figs. 9.11 and 9.12 show the scatter plots and Bland-Altman plots of the automated results to the manually obtained measurements, respectively. The Pearson's correlation coefficients obtained using Approaches I, II and the first and last iteration of Approach III were found to be 0.931, 0.950, 0.975 and 0.984, respectively. Those correlation coefficients were compared using the Fisher's z -transformation and the coefficients obtained using Approaches I and II were found to significantly different ($p < 0.01$) from those obtained using Approach III. However, the coefficients obtained when using the first or last iteration of Approach III did not differ from each other.

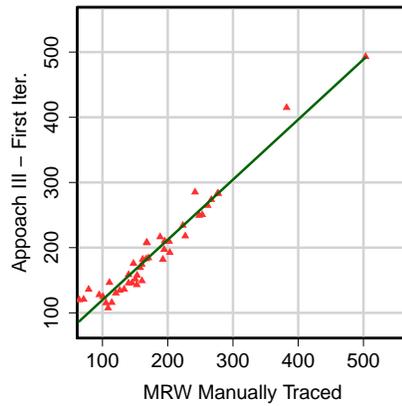
Fig. 9.13(a) shows the iterative results obtained on a radial slice with the first and second iteration results marked in blue and green, respectively. This dataset showed the largest 3-D error, which on the right side is $210.33 \mu\text{m}$. It is evident that the iterative method has found the end of the border tissue rather than the NCO, leading to these large errors. The MRW measurements computed using the manually delineated point and the automated result however, did not differ greatly and were found to be $257.21 \mu\text{m}$ and $245.79 \mu\text{m}$, respectively. Fig. 9.13(b) shows the iterative results obtained on a radial slice with the first, second and third iteration results marked in blue, green and yellow, respectively. This dataset showed the largest MRW difference and it can be seen that the detection of the border tissue along with the shallow cup lead to the large ($\sim 100 \mu\text{m}$) difference in the MRW measurements.



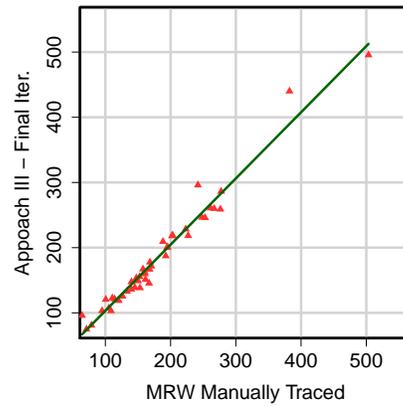
(a)



(b)



(c)



(d)

Figure 9.11: Correlation plots of MRW measurements obtained using manual delineations and the automated methods. The plots for (a) Approach I, (b) Approach II, (c) the first iteration of Approach III and (d) the final iteration of Approach III. The green line shows the linear correlation

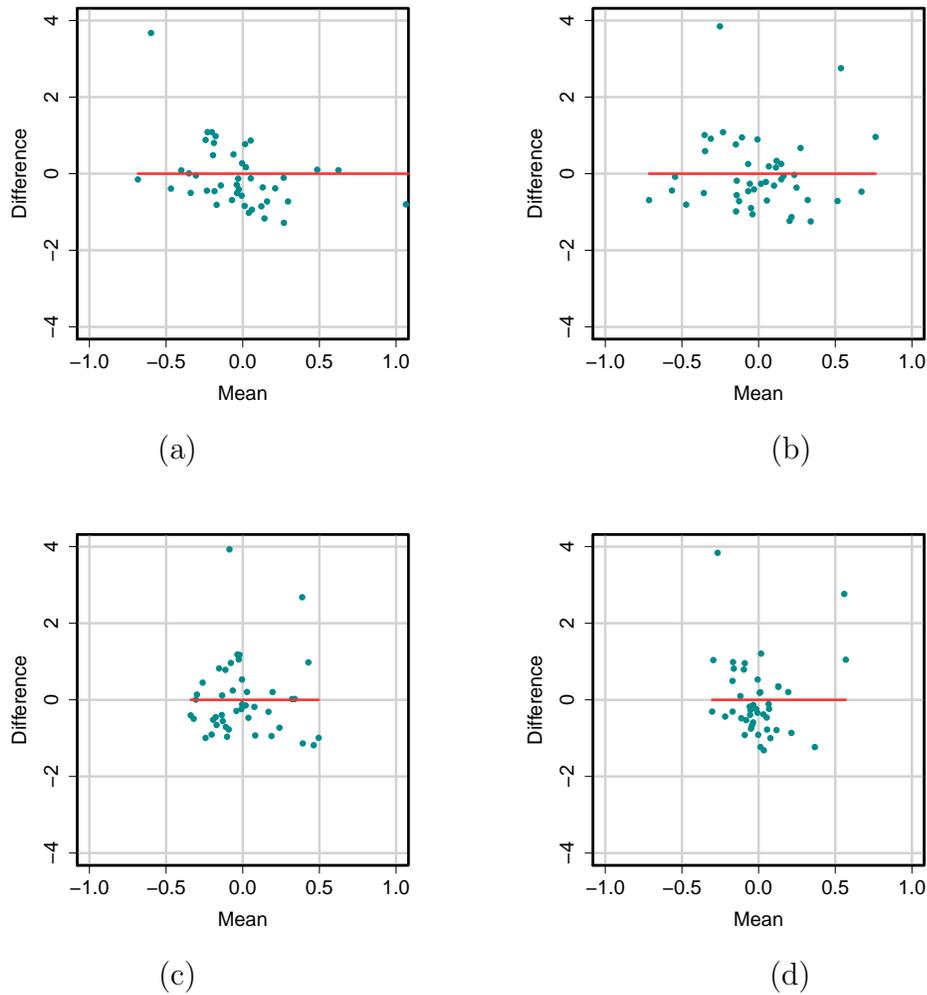
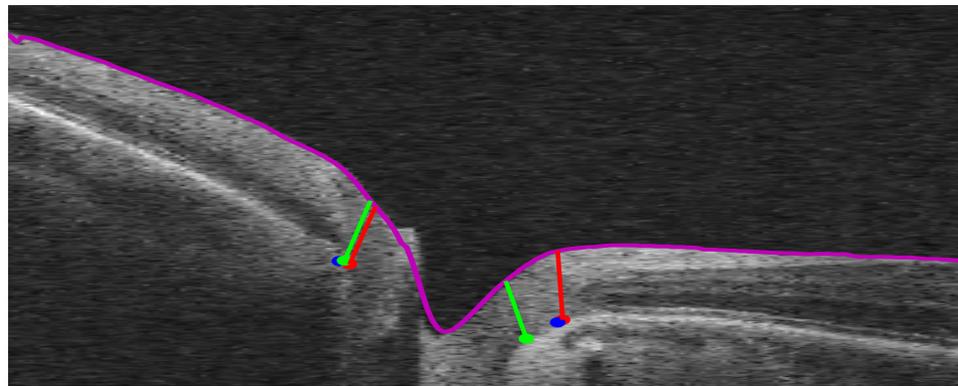
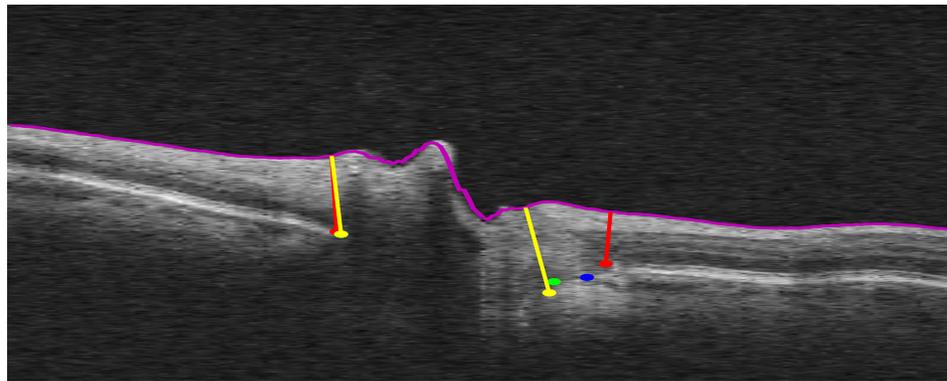


Figure 9.12: Bland-Altman plots of MRW measurements obtained using manual delineations and the automated methods. The plots for (a) Approach I, (b) Approach II, (c) the first iteration of Approach III and (d) the final iteration of Approach III.



(a)



(b)

Figure 9.13: Example of where the iterative method erroneously detected the end of the border tissue instead of the NCO. (a) A slice from a dataset that showed the largest z depth error. (b) A slice from a dataset that showed the largest MRW difference.

9.6 Conclusion & Discussion

We have presented an iterative graph-based method for the automated simultaneous segmentation of surfaces with a shared hole. The graph-theoretic formulation proposed previously [21, 23] was reformulated to meet the new requirements of surfaces that meet at a shared hole boundary, thereby allowing for the boundary of the hole to be detected in 3-D. The proposed method was applied to the segmentation of the NCO in SD-OCT images obtained from human patients that presented with varying stages of glaucoma, where it showed good accuracy in the 2-D plane with errors of 1.54 voxels ($46.05 \pm 28.40 \mu\text{M}$), which is comparable to previously obtained results [28, 100]. It must be noted however, that the results presented by Hu *et al.* [28] compared the 2-D segmentation result to tracings obtained on fundus photographs. The 3-D error noted using the proposed iterative approach was larger ($29.58 \pm 20.24 \mu\text{m}$) than those presented earlier by Hu [100], however it is important to note that the 3-D evaluation presented in that work was conducted using primate data obtained on a Spectralis SD-OCT scanner (where the scans have higher resolution). The border tissue that characterizes the human scans was also not as distinctly visible in these primate scans.

The proposed method also showed improvements from the first iteration to the last iteration, further emphasizing the advantage of using the proposed iterative methods over previously proposed two-step approaches. The MRW measurements obtained using the proposed approach were not significantly different from those obtained using the manually delineated points. Moreover, these measurements correlated closely ($r^2 = 0.98$) with the manually obtained MRW measurements. It is also important to note that the definition of the MRW and its potential use as a predictor of glaucoma is still an area of active research. Reis *et al.* [43, 101], defined the MRW as the shortest perpendicular distance between the NCO and the ILM, thereby accounting for the trajectory of the axons over the point of the measurement. This definition in turn

raises more questions, such as, is this measurement defined individually in each B-scan or is it defined in 3-D? Given the availability of volumetric scans, defining the structure in 3-D would be preferable. In such a situation, it might also be necessary to impose a smoothness constraint similar to the one imposed on the NCO boundary (see Eq. 9.5), as this smooth “cup”-like structure would be similar to the cup studied in planimetry. Furthermore, given that the thickness of the axonal tissue is of interest, perhaps a surface defined from the NCO to the end point of the border tissue would provide for a better definition of the MRW.

The biggest stumbling block for the iterative approach proved to be the presence of the end point of the border tissue, which led to the large errors seen in depth. The machine-learning based “corner” detection faired reasonably well given the task at hand, which was to locate 360 points within a substantially larger volume. However, as the NCO and the endpoint of the border tissue share “corner”-like features, the classifier was not able to differentiate between them, leading to the preferential detection of the border tissue.

Despite the errors that appear substantial in 3-D and along the z -axis, the MRW measurements obtained show that these automated segmentations could still be extremely useful. As shown in Fig. 9.13(a) where the 3-D error was extremely difficult to ignore, but the MRW measurements obtained at the same location was quite similar. Thus, the ability of the proposed method to derive clinically significant rim parameters that could augment planimetry assessments cannot be discounted.

CHAPTER 10

INVESTIGATION OF TEXTURAL MARKERS DERIVED FROM SD-OCT IMAGES THAT MAY PREDICT STRUCTURAL AND FUNCTIONAL CHANGE IN DEGENERATIVE DISEASES

Glaucoma is known to cause degeneration of the retinal ganglion cells (RGC) in the nerve fiber bundles at the retina as well as the ONH, with eventual degeneration of the ganglion cell soma via apoptotic cell death [102]. This degenerative structural change, which occurs predominantly along the nerve fiber bundle tract, was studied quantitatively and was seen to correlate with the peripapillary nerve fiber thickness [37] as well as the thickness of the rim [38]. Zhang *et al.* [103] has also demonstrated that the structural change can be used to derive the visual field loss.

The structural change (thinning) is measurable after cell death occurs. This process occurs over time and the early stages of “sickness” of the cells may affect the reflectance properties of the cells. This may result in a quantitative textural changes in the SD-OCT images. The objective here, is to study retinal texture in images obtained from normal as well as diseased datasets in order to find any statistically significant features that mark disease progression.

10.1 Method

The data used in this experiment consisted of two datasets obtained from human subjects and one obtained from canines. The data obtained from human subjects consisted of SD-OCT images centered on the macula obtained from the right eyes of 45 normal subjects, and 58 subjects that presented with early glaucoma or were suspected of glaucoma, acquired on a Cirrus (Carl Zeiss Meditec, Inc., Dublin CA) scanner. Each dataset has dimensions of $200 \times 200 \times 1024$ voxels obtained from a region $6\text{mm} \times 6\text{mm} \times 2\text{mm}$ volume centered on the macula.

The canine datasets (previously described in Section 7.2.1) were acquired from a colony of Basset hounds, where the canines were divided in three groups: ones that showed genetic susceptibility to glaucoma and presented with the disease (glaucoma

set), a set that showed genetic susceptibility to glaucoma but did not develop the disease (carriers) and a third set of non-affected animals (normals). The baseline scans were obtained at six months of age, and then subsequently prior to 1-4 provocative tests involving the acute elevation of IOP (AEIOP). In this experiment, scans obtained from the right eye at three time points were used: the baseline scans and scans obtained before the first and third AEIOP tests. The carriers were only scanned before the first AEIOP test, and therefore this set only consisted of two time points. The entire study consisted of 24 glaucomatous canines, 13 carriers and 6 normals. However, only 12 glaucomatous canines, 8 carriers and 6 normals had scans acquired from the right eye at the each of the time points of interest. Thus, the final set of images used in this experiment consisted of 12 glaucoma, 8 carriers and 6 nonaffected animals.

10.1.1 Human Data

The images obtained from the human subjects were first segmented to give us the retinal nerve fiber layer (RNFL) and the ganglion cell layer (GCL). The textural markers computed included Gabor energy features computed at scales $\sigma = \{1,2,4\}$ and angles $\theta = \{0^\circ, 30^\circ, 45^\circ, 60^\circ, 90^\circ\}$ and intensity based features (see Section 5.1). The average of the textural markers were then computed within 6 sectoral regions for the RNFL (see Figs. 10.1(a) - (c)) and 7 regions for the GCL as shown in Figs. 10.1(d) and (e). As the thinning seen in the RNFL associated with glaucoma is known to occur along the main arcuate, the masks were designed to encompass these regions. Similarly, the thinning seen in the GCL typically occurs in an annular pattern which is reflected in the shape of the masks used.

The mean textural feature derived from within each sector from the normals, glaucoma suspects and glaucoma patients were then compared statistically using t -tests (corrected for multiple corrections) in order to gauge if any significant differences exist among these groups. The thickness of the RNFL and the GCL was also computed

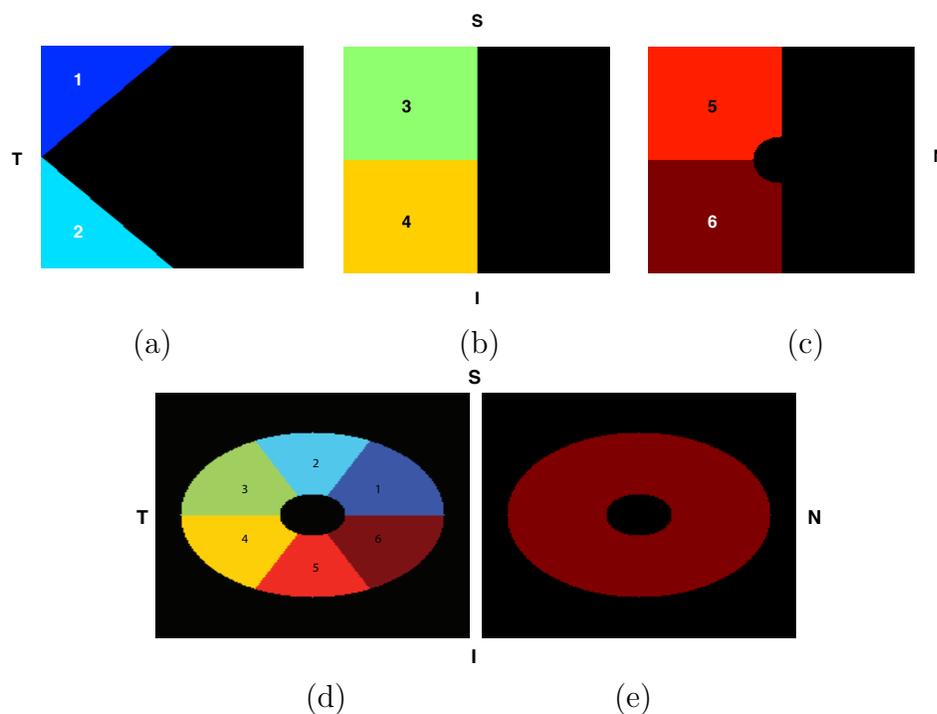


Figure 10.1: Sector maps used to analyze human scans. Panels (a), (b) and (c) show the sectors of interest that were used to investigate texture in the RNFL. Panels (d) and (e) how the seven sectors of interest that were used to investigate texture in the GCL.

within these regions and statistically compared to each other.

10.1.2 Canine Data

The RNFL and the GCL in canines appears as a combined layer (as described previously in Section 7.2.1) that we refer to as the NF+GC complex. This layer can be quite thin in some regions of the retina and the reproducibility of thickness measurements associated with this layer can be lower than for the structure comprised of the NF+GC complex and IPL (as presented in Section 6.3). Thus, the layer used in this experiment consisted of the NF+GC complex and the IPL. The textural features described above were generated for these scans and the mean was computed using the mask shown in Fig. 10.2. Recall from Section 7.2.1, these scans typically contained

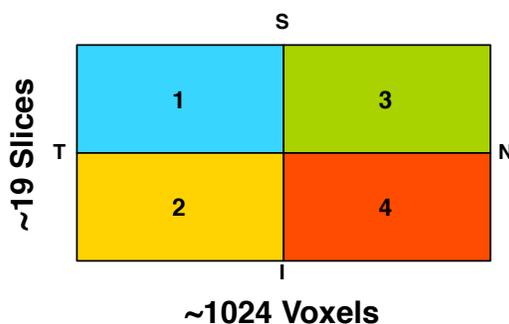


Figure 10.2: The sector map used to analyze the dog scans. The illustration shows the 4 quadrants within which the texture of the NF+GC+IPL was computed.

19 slices and thus the use of a more detailed mask was not possible. Although these scans were macula-centered, the ONH is visible in these scans and occurred within the third and fourth quadrants of the mask.

The mean textural feature derived from within these four sectors from the normals, carriers and glaucoma canines. The mean thickness of the NF+GC+IPL was also computed within these sectors. Since the potential textural change over time is of interest, the mean thickness and texture at the 3 time points within each group of animals were compared statistically using *t*-tests (corrected for multiple corrections). The mean thickness and textural features obtained from the three groups were not compared to each other.

10.2 Results

10.2.1 Human Data

The mean thickness of the RNFL and the GCL observed in the sectors are presented in Table 10.1. The mean thickness of the RNFL in glaucoma subjects in sectors 2 and 4 were seen to be significantly smaller ($p < 0.001$) than the RNFL thickness observed in the glaucoma suspect group and the normals. The RNFL thickness within those sectors noted in the normals and the suspects did not significantly differ from each other. However, a subset of the texture features (mean intensity, variance and

Table 10.1: The mean thickness (in mean \pm standard deviation μm) of the RNFL and the GCL computed within the predefined sectors.

	Sector	Normals	Suspects	Glaucoma
RNFL	1	61.29 \pm 13.77	61.58 \pm 13.74	49.02 \pm 16.01
	2	61.31 \pm 10.76	60.97 \pm 19.10	43.45 \pm 16.23
	3	46.71 \pm 9.72	46.82 \pm 9.86	39.64 \pm 10.82
	4	45.74 \pm 7.06	46.05 \pm 12.70	35.49 \pm 11.01
	5	47.88 \pm 9.95	47.96 \pm 10.11	40.57 \pm 11.22
	6	46.94 \pm 7.29	47.19 \pm 13.16	36.32 \pm 11.45
GCL	1	79.85 \pm 8.90	75.75 \pm 12.33	65.95 \pm 10.58
	2	76.76 \pm 10.30	74.42 \pm 11.12	64.24 \pm 12.27
	3	79.70 \pm 11.66	77.01 \pm 10.51	67.80 \pm 11.68
	4	79.13 \pm 10.00	74.98 \pm 12.11	69.05 \pm 11.50
	5	81.52 \pm 8.92	76.16 \pm 15.13	68.13 \pm 11.99
	6	81.90 \pm 8.92	78.37 \pm 15.45	65.81 \pm 12.24
	7	80.53 \pm 11.66	74.32 \pm 12.59	60.95 \pm 12.65

entropy) within these two sectors did significantly differ between the normals and the suspects ($p < 0.001$), while the same features computed within the glaucoma set did not significantly differ from the suspects.

Fig. 10.3 shows the bar plot of 3 features that showed significant differences among the normals and the suspects in sectors 2 and 4 of the RNFL. Most sectors (all except sector 2) in the GCL showed significant thickness change among the suspects and the glaucoma patients, with no significant difference between the normals and the suspects. However, only sector 5 showed a significant textural difference among the normals and the suspects. The same sector showed no difference among the suspects and the glaucoma patients.

10.2.2 Canine Data

The mean thickness of the combined NG+GC complex and the IPL in the canine data in the 4 quadrants are summarized in Table 10.2. As mentioned previously, the carriers were not imaged beyond the first AEIOP session and thus, the third time

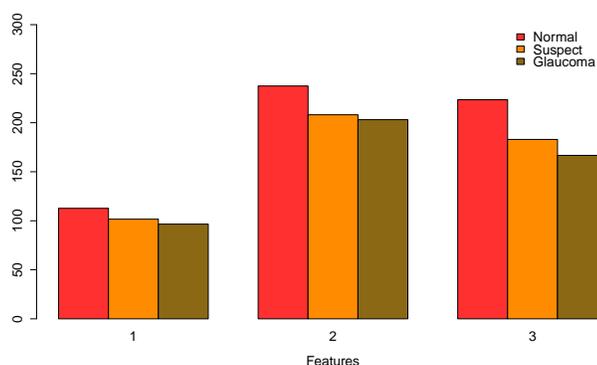


Figure 10.3: Bar plot of features obtained from normals, suspects and glaucoma patients. The plot shows the mean feature value of three intensity-based features derived within sector 3 for the normal, suspect and glaucoma datasets.

point data was not available. The thickness measurements obtained on the normals and the carriers did not differ significantly at any of the time points. In the glaucoma set, however, a significant reduction ($p < 0.01$) was noted in quadrants 1, 2 and 3 in the thickness measurements obtained at baseline and at the first time point (before the first AEIOP). The thickness within these same sectors obtained at baseline were also significantly ($p < 0.01$) thicker than those obtained at the third time point (before the third AEIOP). No significant differences were noted in the texture in the glaucoma, carrier or normal set at any time point within any of the sectors.

10.3 Discussion & Conclusion

To investigate the hypothesis that textural change may precede structural RGC death, textural features were computed within three sets of data. In the human set, three intensity-based features showed a significant difference between the normals and the suspects, while the same features did not significantly differ among the glaucoma suspects and the glaucoma patients. While this is an encouraging result, the difference could very well be correlated with age rather than disease as the normals were not age-matched with the suspects and glaucoma patients.

Table 10.2: Mean thickness of NF+GC+IPL in predefined sectors in canine data. The mean thickness (in mean \pm standard deviation μm) was computed within the 4 predefined quadrants in the canine data.

	T.P.	1	2	3	4
Norm.	1	97.11 \pm 8.86	139.58 \pm 27.67	91.86 \pm 8.00	128.82 \pm 15.90
	2	93.82 \pm 6.55	141.57 \pm 27.29	101.97 \pm 25.88	128.16 \pm 26.30
	3	93.58 \pm 7.06	138.12 \pm 17.70	101.66 \pm 23.46	130.97 \pm 20.07
Car.	1	95.23 \pm 2.55	154.23 \pm 15.26	93.35 \pm 5.07	135.62 \pm 12.80
	2	96.82 \pm 4.00	155.19 \pm 18.09	92.33 \pm 5.10	134.01 \pm 10.51
	3	-	-	-	-
Glau.	1	84.64 \pm 4.67	132.35 \pm 15.70	81.75 \pm 4.65	115.04 \pm 12.93
	2	81.64 \pm 4.75	122.33 \pm 16.80	78.15 \pm 10.05	117.12 \pm 12.84
	3	81.50 \pm 4.22	121.17 \pm 13.05	79.90 \pm 5.62	116.35 \pm 18.46

The canine dataset was a good alternative as the study contained a well defined control set. However, the data consisted of sparsely acquired volumetric scans where the images contained 13 to 23 slices. Furthermore, the glaucomatous set showed a significant thinning in the NF+GC and IPL layer between time points 1 and 2. Ideally, we would like to have a set of images acquired before the onset of structural change that would allow us to test the proposed hypothesis.

The lack of any textural change even after the on-set of glaucomatous change is a slightly discouraging result. However, the original study protocol called for the treatment of the canines with topical and oral drugs that at all 3 time points. The drugs were expected to reverse effects of the AEIOP [104] and thus, could have effected the images acquired from these animals. In the future, we would like to repeat this experiment while studying the progression of a degenerative retinal disease in an animal model, without external treatment plans or AEIOP in the animals.

CHAPTER 11 CONCLUSION

The overarching theme of the work presented here has been the incorporation of *a priori* information into automated methods in order to improve their performance. In Chapter 4, a method for the correction of the characteristic axial artifacts frequently seen in SD-OC images was described. The predictable nature of the artifacts was leveraged here, allowing for them to be better modeled with two splines rather than one.

In Chapters 5 and 7, a combined machine-learning and graph-theoretic approach was presented, where terms of the cost function were automatically designed by learning texture features from a training set. This not only significantly improved the segmentation accuracy of the method, it also made the method easily adaptable to new models, such as mice and canine scans. However, the impact of the learned features need to be gauged carefully as the learned components may not always improve the accuracy. This was observed to be the case for the glaucomatous scans obtained from human patients and canines, where the learned in-region cost term significantly improved the segmentation accuracy but addition of the learned on-surface cost term did not. Moreover, certain surfaces may prefer either the learned on-surface cost or the learned in-region cost rather than the combination of both. Thus, the final graph-theoretic method could include customizations for each surface as well. However, given the complexity of the interactions that exist when segmentation multiple surfaces make this a difficult optimization problem to solve and thus, was not addressed in this work.

The incorporation of learned terms into the cost function was extended further in Chapter 8, where learned shape priors were incorporated into the graph-theoretic method. While the inclusion of the learned shape priors improved the segmentation accuracy with respect to the baseline results, its inclusion into a system that already

incorporated learned in-region and on-surface cost terms did not raise the segmentation accuracy. Thus, we can conclude that the incorporation of learned features can enhance the performance of the graph-theoretic approach, but the blind incorporation of learned components may not provide the best results.

In Chapter 9, we described a method for the automated simultaneous segmentation of surfaces and their shared hole. Here, we saw that the suitable adjustments made to the constraints and the cost function utilized by the graph-theoretic approach significantly improved the segmentation accuracy as compared to previously described two-step approaches. In addition to being able to detect the NCO, important measurements such as the MRW can also be obtained using this approach. In particular, the close correlation between the MRW measurements obtained using the automated method and the measurements obtained using the manually delineated points indicate that the errors noted in 3-D, while large in some cases, may not be crucial to the application. The definition of MRW is also, at the moment, not well defined. Further investigations into the relationship of the rim width at the NCO and the disease will provide a more appropriate definition of this measurement, and we expect that this automated method will help in the detection of the structures of interest.

In Chapter 10, the hypothesis that textural markers derived from SD-OCT images could be used to predict structural change seen in degenerative diseases (such as glaucoma) was investigated. While the results did not give us definitive results, it demonstrates the potential uses of automated segmentation approaches without which, such a study would not have been possible without considerable manual intervention.

The applications presented here were focused on the segmentation and quantitative study of retinal structures in SD-OCT images. However, the methods could be adapted to other image analysis problems (where multiple layers structures need

to be segmented) as well, as the main requirement of the methods presented here is an appropriate training set. In conclusion, the combining of machine learning-based methods and a graph-theoretic approach not only provides improvements in segmentation accuracy, but also makes the methods easily adapted to new applications.

REFERENCES

- [1] D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, and C. A. Puliafito, "Optical coherence tomography," *Science*, vol. 254, no. 5035, pp. 1178–1181, Nov. 1991.
- [2] D. Cabrera Fernández, H. M. Salinas, and C. A. Puliafito, "Automated detection of retinal layer structures on optical coherence tomography images," *Opt. Express*, vol. 13, no. 25, pp. 10 200–10 216, 2005.
- [3] M. Baroni, P. Fortunato, and A. L. Torre, "Towards quantitative analysis of retinal features in optical coherence tomography," *Med Eng Phys*, vol. 29, no. 4, pp. 432–441, May 2007.
- [4] M. Shahidi, Z. Wang, and R. Zelkha, "Quantitative thickness measurement of retinal layers imaged by optical coherence tomography," *Am. J. Ophthalmol.*, vol. 139, no. 6, pp. 1056–1061, Jun. 2005.
- [5] A. Chan, J. S. Duker, H. Ishikawa, T. H. Ko, J. S. Schuman, and J. G. Fujimoto, "Quantification of photoreceptor layer thickness in normal eyes using optical coherence tomography." *Retina*, vol. 26, no. 6, pp. 655–660, 2006.
- [6] H. Ishikawa, D. M. Stein, G. Wollstein, S. Beaton, J. G. Fujimoto, and J. S. Schuman, "Macular segmentation with optical coherence tomography," *Invest. Ophthalmol. Vis. Sci.*, vol. 46, no. 6, pp. 2012–2017, Jun. 2005.
- [7] A. M. Bagci, M. Shahidi, R. Ansari, M. Blair, N. P. Blair, and R. Zelkha, "Thickness Profiles of Retinal Layers by Optical Coherence Tomography Image Segmentation," *Am. J. Ophthalmol.*, vol. 146, no. 5, pp. 679–687, Nov. 2008.
- [8] T. Fabritius, S. Makita, M. Miura, R. Myllylä, and Y. Yasuno, "Automated segmentation of the macula by optical coherence tomography." *Opt. Express*, vol. 17, no. 18, pp. 15 659–69, Aug. 2009.
- [9] O. Tan, G. Li, A. T.-H. Lu, R. Varma, D. Huang, and Advanced Imaging for Glaucoma Study Group, "Mapping of macular substructures with optical coherence tomography for glaucoma diagnosis," *Ophthalmology*, vol. 115, no. 6, pp. 949–956, Jun. 2008.
- [10] O. Tan, V. Chopra, A. T.-H. Lu, J. S. Schuman, H. Ishikawa, G. Wollstein, R. Varma, and D. Huang, "Detection of macular ganglion cell loss in glaucoma by Fourier-domain optical coherence tomography," *Ophthalmology*, vol. 116, no. 12, pp. 2305–2314, Dec. 2009.

- [11] R. J. Zawadzki, A. R. R. Fuller, D. F. F. Wiley, B. Hamann, S. S. S. Choi, and J. S. S. Werner, "Adaptation of a support vector machine algorithm for segmentation and visualization of retinal structures in volumetric optical coherence tomography data sets," *J. Biomed. Opt.*, vol. 12, no. 4, p. 41206, 2007. [Online]. Available: <http://dx.doi.org/10.1117/1.2772658>
- [12] F. Rossant, I. Ghorbel, I. Bloch, M. Paques, and S. Tick, "Automated segmentation of retinal layers in OCT imaging and derived ophthalmic measures," in *Proc. of IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, 2009, pp. 1370–1373.
- [13] K. A. Vermeer, J. van der Schoot, H. G. Lemij, and J. F. de Boer, "Segmentation of retinal layers in volumetric OCT scans of normal and glaucomatous subjects," in *SPIE Photonics West, Ophthalmic Technologies XXI*, F. Manns, P. G. Soderberg, and A. Ho, Eds., vol. 7885, no. 1. SPIE, 2011, p. 78851B.
- [14] A. Mishra, A. Wong, K. Bizheva, and D. A. Clausi, "Intra-retinal layer segmentation in optical coherence tomography images," *Opt. Express*, vol. 17, no. 26, pp. 23 719–23 728, 2009.
- [15] A. Yazdanpanah, G. Hamarneh, B. R. Smith, and M. V. Sarunic, "Segmentation of intra-retinal layers from optical coherence tomography images using an active contour approach." *IEEE Trans. Med. Imag.*, vol. 30, no. 2, pp. 484–96, Feb. 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20952331>
- [16] F. Rathke, S. Schmidt, and C. Schnörr, "Order preserving and shape prior constrained intra-retinal layer segmentation in optical coherence tomography," in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, vol. 14, Jan. 2011, pp. 370–377.
- [17] V. Kajić, B. Povazay, B. Hermann, B. Hofer, D. Marshall, P. L. Rosin, and W. Drexler, "Robust segmentation of intraretinal layers in the normal human fovea using a novel statistical model based on texture and shape analysis." *Opt. Express*, vol. 18, no. 14, pp. 14 730–44, Jul. 2010.
- [18] Q. Yang, C. A. Reisman, Z. Wang, Y. Fukuma, M. Hangai, N. Yoshimura, A. Tomidokoro, M. Araie, A. S. Raza, D. C. Hood, and K. Chan, "Automated layer segmentation of macular OCT images using dual-scale gradient information," *Opt. Express*, vol. 18, no. 20, pp. 21 293–21 307, 2010.
- [19] S. J. Chiu, X. T. Li, P. Nicholas, C. A. Toth, J. A. Izatt, and S. Farsiu, "Automatic segmentation of seven retinal layers in SDOCT images congruent with expert manual segmentation." *Opt. Express*, vol. 18, no. 18, pp. 19 413–28, Aug. 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20940837>

- [20] S. J. Chiu, J. a. Izatt, R. V. O’Connell, K. P. Winter, C. a. Toth, and S. Farsiu, “Validated automatic segmentation of AMD pathology including drusen and geographic atrophy in SD-OCT images.” *Invest. Ophthalmol. Vis. Sci.*, vol. 53, no. 1, pp. 53–61, Jan. 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22039246>
- [21] M. K. Garvin, M. D. Abràmoff, X. Wu, S. R. Russell, T. L. Burns, and M. Sonka, “Automated 3-D intraretinal layer segmentation of macular spectral-domain optical coherence tomography images,” *IEEE Trans. Med. Imag.*, vol. 28, no. 9, pp. 1436–1447, 2009.
- [22] X. Wu and D. Z. Chen, “Optimal Net Surface Problems with Applications,” in *Proc. of the 29th International Colloquium on Automata, Languages, and Programming (ICALP), LNCS 2380*. Springer-Verlag, 2002, pp. 1029–1042.
- [23] K. Li, X. Wu, D. Z. Chen, and M. Sonka, “Optimal surface segmentation in volumetric images—a graph-theoretic approach.” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 1, pp. 119–134, Jan. 2006.
- [24] M. Haeker, X. X. Wu, M. D. Abràmoff, R. Kardon, and M. Sonka, “Incorporation of regional information in optimal 3- $\{D\}$ graph search with application for intraretinal layer segmentation of optical coherence tomography images,” in *Proceedings of the 20th International Conference on Information Processing in Medical Imaging (IPMI 2007)*, ser. Lecture Notes in Computer Science, vol. 4584, Springer-Verlag. Springer, 2007, pp. 607–618.
- [25] K. Lee, M. Niemeijer, M. K. Garvin, Y. H. Kwon, M. Sonka, and M. D. Abràmoff, “Segmentation of the Optic Disc in 3D-OCT Scans of the Optic Nerve Head,” *IEEE Trans. Image Process.*, vol. 29, no. 1, pp. 159–68, 2009.
- [26] Q. Song, X. Wu, and Y. Liu, “Simultaneous searching of globally optimal interacting surfaces with shape priors,” *IEEE Conference on Computer Vision and Patter Recognition (CVPR 2010)*, pp. 2879–2886, 2010. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5540025
- [27] B. J. Antony, M. D. Abramoff, K. Lee, P. Sonkova, P. Gupta, Y. Kwon, M. Niemeijer, Z. Hu, and M. K. Garvin, “Automated 3D segmentation of intraretinal layers from optic nerve head optical coherence tomography images,” *Proc. of SPIE 2010*, vol. 7626, p. 76260U, 2010. [Online]. Available: <http://link.aip.org/link/PSISDG/v7626/i1/p76260U/s1&Agg=doi>
- [28] Z. Hu, M. D. Abràmoff, Y. H. Kwon, K. Lee, and M. K. Garvin, “Automated Segmentation of Neural Canal Opening and Optic Cup in 3D Spectral Optical

- Coherence Tomography Volumes of the Optic Nerve Head,” *Invest. Ophthalmol. Vis. Sci.*, vol. 51, no. 11, pp. 5708–5717, 2010.
- [29] Y. Yin, X. Zhang, R. Williams, X. Wu, D. D. Anderson, and M. Sonka, “LOGISMOS: Layered Optimal Graph Image Segmentation of Multiple Objects and Surfaces: Cartilage Segmentation in the Knee Joint,” *IEEE Trans. Med. Imag.*, vol. 29, no. 12, pp. 2023–2037, 2010.
- [30] Q. Song, J. Bai, M. K. Garvin, M. Sonka, J. M. Buatti, and X. Wu, “Optimal multiple surface segmentation with shape and context priors,” *IEEE Trans. Med. Imag.*, vol. 32, no. 2, pp. 376–86, Feb. 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23193309>
- [31] B. E. Klein, R. Klein, W. E. Sponsel, T. Franke, L. B. Cantor, J. Martone, and M. J. Menage, “Prevalence of glaucoma. The Beaver Dam Eye Study.” *Ophthalmology*, vol. 99, no. 10, p. 1499, 1992.
- [32] D. S. Friedman, R. C. Wolfs, B. J. O’colmain, B. E. Klein, H. R. Taylor, S. West, M. C. Leske, P. Mitchell, N. Congdon, and J. Kempen, “Prevalence of open-angle glaucoma among adults in the United States,” *Arch. Ophthalmol.*, vol. 122, no. 4, p. 532, 2004.
- [33] J. E. DeLeón-Ortega, S. N. Arthur, G. McGwin, A. Xie, B. E. Monheit, and C. A. Girkin, “Discrimination between glaucomatous and nonglaucomatous eyes using quantitative imaging devices and subjective optic nerve head assessment,” *Invest Ophthalmol Vis Sci*, vol. 47, no. 8, pp. 3374–3380, Aug. 2006.
- [34] H.-Y. Chen and M.-L. Huang, “Discrimination between normal and glaucomatous eyes using Stratus optical coherence tomography in Taiwan Chinese subjects,” *Graefes. Arch. Clin. Exp. Ophthalmol.*, vol. 243, no. 9, pp. 894–902, 2005. [Online]. Available: <http://dx.doi.org/10.1007/s00417-005-1140-y>
- [35] F. Medeiros, L. Zangwill, C. Bowd, R. Vessani, S. Remo Jr., and R. N. Weinreb, “Evaluation of Retinal Nerve Fiber Layer, Optic Nerve Head, and Macular Thickness Measurements for Glaucoma Detection Using Optical Coherence Tomography,” *Am. J. Ophthalmol.*, vol. 139, no. 1, pp. 44–55, 2005.
- [36] G. Li, A. K. Fansi, J.-F. Boivin, L. Joseph, and P. Harasymowycz, “Screening for glaucoma in high-risk populations using optical coherence tomography.” *Ophthalmology*, vol. 117, no. 3, pp. 453–61, Mar. 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20031231>

- [37] M. K. Garvin, M. D. Abramoff, K. Lee, M. Niemeijer, M. Sonka, and Y. H. Kwon, "2-D Pattern of Nerve Fiber Bundles in Glaucoma Emerging From Spectral-Domain Optical Coherence Tomography." *Invest. Ophthalmol. Vis. Sci.*, vol. 53, no. 1, pp. 483–9, Jan. 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22222272>
- [38] K. Lee, Y. H. Kwon, M. K. Garvin, M. Niemeijer, M. Sonka, and M. D. Abramoff, "Distribution of Damage to the Entire Retinal Ganglion Cell Pathway," *Arch. Ophthalmol.*, p. 1, May 2012. [Online]. Available: <http://archophth.jamanetwork.com/article.aspx?doi=10.1001/archophthalmol.2012.669>
- [39] A. J. Bellezza, C. J. Rintalan, H. W. Thompon, J. C. Downs, R. T. Hart, and C. F. Burgoyne, "Deformation of the Lamina Cribrosa and Anterior Scleral Canal Wall in Early Experimental Glaucoma," *Invest. Ophthalmol. Vis. Sci.*, vol. 44, no. 2, pp. 623–637, Feb. 2003. [Online]. Available: <http://www.iovs.org/cgi/doi/10.1167/iovs.01-1282>
- [40] H. Yang, J. C. Downs, A. Bellezza, H. Thompson, and C. F. Burgoyne, "3-D histomorphometry of the normal and early glaucomatous monkey optic nerve head: prelaminar neural tissues and cupping." *Invest. Ophthalmol. Vis. Sci.*, vol. 48, no. 11, pp. 5068–84, Nov. 2007. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17962459>
- [41] J. C. Downs, H. Yang, C. Girkin, L. Sakata, A. Bellezza, H. Thompson, and C. F. Burgoyne, "Three-dimensional histomorphometry of the normal and early glaucomatous monkey optic nerve head: neural canal and subarachnoid space architecture." *Invest Ophthalmol Vis Sci*, vol. 48, no. 7, pp. 3195–3208, Jul. 2007. [Online]. Available: <http://dx.doi.org/10.1167/iovs.07-0021>
- [42] N. G. Strouthidis, H. Yang, J. C. Downs, and C. F. Burgoyne, "Comparison of clinical and three-dimensional histomorphometric optic disc margin anatomy." *Invest. Ophthalmol. Vis. Sci.*, vol. 50, no. 5, pp. 2165–74, May 2009.
- [43] A. S. C. Reis, N. O'Leary, H. Yang, G. P. Sharpe, M. T. Nicolela, C. F. Burgoyne, and B. C. Chauhan, "Influence of clinically invisible, but optical coherence tomography detected, optic disc margin anatomy on neuroretinal rim evaluation." *Invest. Ophthalmol. Vis. Sci.*, vol. 53, no. 4, pp. 1852–60, Apr. 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22410561>
- [44] N. G. Strouthidis, H. Yang, B. Fortune, J. C. Downs, and C. F. Burgoyne, "Detection of the optic nerve head neural canal opening within three-dimensional histomorphometric and spectral domain optical coherence tomography data sets," *Invest. Ophthalmol. Vis. Sci.*, vol. 50, no. 1, p. 214, 2009.

- [45] M. Haeker, M. D. Abràmoff, X. Wu, R. Kardon, and M. Sonka, "Use of varying constraints in optimal 3-D graph search for segmentation of macular optical coherence tomography images," in *Proc. of Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI 2007)*, ser. Lecture Notes in Computer Science, vol. 4791. Springer-Verlag, 2007, pp. 244–251. [Online]. Available: <http://www.springerlink.com/index/y5j0w10544321181.pdf>
- [46] Q. Song, X. Wu, Y. Liu, M. Smith, J. Buatti, and M. Sonka, "Optimal graph search segmentation using arc-weighted graph for simultaneous surface detection of bladder and prostate." *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, vol. 12, no. Pt 2, pp. 827–35, Jan. 2009. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20426188>
- [47] D. Koozekanani, K. Boyer, and C. Roberts, "Retinal thickness measurements from optical coherence tomography using a Markov boundary model," *IEEE Trans. Med. Imag.*, vol. 20, no. 9, pp. 900–916, 2001.
- [48] X. Zhang, S. Yousefi, L. An, and R. K. Wang, "Automated segmentation of intramacular layers in Fourier domain optical coherence tomography structural images from normal subjects." *J. Biomed. Opt.*, vol. 17, no. 4, p. 046011, Apr. 2012.
- [49] S. Lu, C. Y.-I. Cheung, J. Liu, J. H. Lim, C. K.-s. Leung, and T. Y. Wong, "Automated layer segmentation of optical coherence tomography images." *IEEE Trans. Biomed. Eng.*, vol. 57, no. 10, pp. 2605–8, Oct. 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20595078>
- [50] A. Lang, A. Carass, E. Sotirchos, P. Calabresi, and J. L. Prince, "Segmentation of retinal OCT images using a random forest classifier," in *Proc. SPIE 2013*, vol. 8669, Mar. 2013, p. 86690R.
- [51] M. Mujat, R. Chan, B. Cense, B. Park, C. Joo, T. Akkin, T. Chen, and J. de Boer, "Retinal nerve fiber layer thickness map determined from optical coherence tomography images." *Opt. Express*, vol. 13, no. 23, pp. 9480–91, Nov. 2005. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/19503151>
- [52] X. Xu, M. Niemeijer, Q. Song, M. Sonka, M. K. Garvin, J. M. Reinhardt, and M. D. Abràmoff, "Vessel boundary delineation on fundus images using graph-based approach," *IEEE Transactions on Medical Imaging*, vol. 30, no. 6, pp. 1184–1191, Jun. 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/21216707>
- [53] Q. Song, M. Chen, J. Bai, M. Sonka, and X. Wu, "Surface-region context in optimal multi-object graph-based segmentation: robust delineation of pulmonary

- tumors.” *Information Processing in Medical Imaging (IPMI)*, vol. 22, pp. 61–72, Jan. 2011.
- [54] Q. Song, X. Wu, X. Dou, and M. Sonka, “Globally optimal 3D graph search incorporating both edge and regional information: application to aortic MR image segmentation,” *Proc. of SPIE Medical Imaging: Image Processing (SPIE Medical Imaging 2009)*, vol. 7259, 2009. [Online]. Available: <http://144.206.159.178/FT/CONF/16428378/16428417.pdf> <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1335539>
- [55] A. Lang, A. Carass, M. Hauser, E. S. Sotirchos, P. A. Calabresi, H. S. Ying, and J. L. Prince, “Retinal layer segmentation of macular OCT images using boundary classification,” *Biomed. Opt. Express*, vol. 4, no. 7, p. 1133, Jun. 2013. [Online]. Available: <http://www.opticsinfobase.org/abstract.cfm?URI=boe-4-7-1133>
- [56] L. Breiman, “Random forests,” *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001. [Online]. Available: <http://www.springerlink.com/index/U0P06167N6173512.pdf>
- [57] P. A. Dufour, L. Ceklic, H. Abdillahi, S. Schröder, S. De Dzanet, U. Wolf-Schnurrbusch, and J. Kowal, “Graph-based multi-surface segmentation of OCT data using trained hard and soft constraints,” *IEEE Trans. Med. Imag.*, vol. 32, no. 3, pp. 531–543, Mar. 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23086520>
- [58] M. Wojtkowski, V. Srinivasan, J. G. Fujimoto, T. Ko, J. S. Schuman, A. Kowalczyk, and J. S. Duker, “Three-dimensional retinal imaging with high-speed ultrahigh-resolution optical coherence tomography.” *Ophthalmology*, vol. 112, no. 10, pp. 1734–46, Oct. 2005.
- [59] M. Niemeijer, M. K. Garvin, K. Lee, B. van Ginneken, M. D. Abramoff, and M. Sonka, “Registration of 3D spectral OCT volumes using 3D SIFT feature point matching,” in *Proc. of SPIE Medical Imaging: Image Processing (SPIE Medical Imaging 2009)*, vol. 7259, 2009, p. 72591I.
- [60] P. Thévenaz, U. E. Ruttimann, M. Unser, and P. Thevenaz, “A pyramid approach to subpixel registration based on intensity,” *IEEE Trans. Image Process.*, vol. 7, no. 1, pp. 27–41, Jan. 1998. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18267377>
- [61] M. D. Abramoff, P. J. Magalhaes, and S. J. Ram, “Image processing with ImageJ,” *Biophotonics International*, vol. 11, no. 7, pp. 36–42, 2004.

- [62] R. J. Zawadzki, A. R. Fuller, S. S. Choi, D. F. Wiley, B. Hamann, and J. S. Werner, "Correction of motion artifacts and scanning beam distortions in 3D ophthalmic optical coherence tomography imaging," in *SPIE*, vol. 6426, 2007, pp. 642 601–642 607.
- [63] A. Fuller, R. Zawadzki, S. Choi, D. Wiley, J. Werner, and B. Hamann, "Segmentation of three-dimensional retinal image data," *IEEE Trans. Vis. Comput. Graphics*, vol. 13, no. 6, pp. 1719–1726, 2007.
- [64] A. A. Khanifar, A. F. Koreishi, J. A. Izatt, and C. A. Toth, "Drusen ultrastructure imaging with spectral domain optical coherence tomography in age-related macular degeneration," *Ophthalmology*, vol. 115, no. 11, pp. 1883–1890, 2008.
- [65] S. Ricco, M. Chen, H. Ishikawa, G. Wollstein, and J. Schuman, "Correcting Motion Artifacts in Retinal Spectral Domain Optical Coherence Tomography via Image Registration," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, vol. 5761/2009. Springer, 2009, pp. 100–107.
- [66] M. Kraus, M. A. Mayer, R. Bock, B. Potsaid, V. Manjunath, J. S. Duker, J. Hornegger, and J. G. Fujimoto, "Motion Artifact Correction in OCT Volume Scans Using Image Registration," *Invest. Ophthalmol. Vis. Sci. E-Abstract*, vol. 51, p. 4405, 2010.
- [67] X. Song, R. Estrada, S. J. Chiu, A.-H. Dhalla, C. A. Toth, J. A. Izatt, and Sina Farsiu, "Segmentation-based registration of retinal optical coherence tomography images with pathology," *Invest. Ophthalmol. Vis. Sci. E-Abstract*, vol. 52, p. 1309, 2011.
- [68] A. Hofman, M. M. B. Breteler, C. M. van Duijn, H. L. a. Janssen, G. P. Krestin, E. J. Kuipers, B. H. C. Stricker, H. Tiemeier, A. G. Uitterlinden, J. R. Vingerling, and J. C. M. Witteman, "The Rotterdam Study: 2010 objectives and design update," *European Journal of Epidemiology*, vol. 24, no. 9, pp. 553–72, Jan. 2009.
- [69] R. C. W. Wolfs, P. H. Borger, R. S. Ramrattan, C. C. W. Klaver, C. A. A. Hulsman, A. Hofman, J. R. Vingerling, R. A. Hitchings, and P. T. V. M. de Jong, "Changing views on open-angle glaucoma: Definitions and prevalences-the Rotterdam Study," *Invest Ophthalmol Vis Sci*, vol. 41, pp. 3309–3321, 2000.
- [70] M. A. Czudowska, W. D. Ramdas, and R. C. W. Wolfs, "Incidence of glaucomatous visual field loss : A ten-year follow-up from the Rotterdam study," *Ophthalmology*, vol. 117, no. 9, pp. 1705–1712, 2010.

- [71] J. Duchon, “Splines minimizing rotation-invariant semi-norms in Sobolev spaces,” *Constructive theory of functions of several variables*, pp. 85–100, 1977.
- [72] F. Bookstein, “Principal warps: thin-plate splines and the decomposition of deformations,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 6, pp. 567–585, Jun. 1989.
- [73] M. J. D. Powell, “The uniform convergence of thin plate spline interpolation in two dimensions,” *Numerische Mathematik*, vol. 68, no. 1, pp. 107–128, 1994.
- [74] M. Niemeijer, M. K. Garvin, B. van Ginneken, M. Sonka, and M. D. Abramoff, “Vessel Segmentation in 3-D Spectral OCT Scans of the Retina,” in *Proc. of SPIE Medical Imaging: Image Processing (SPIE Medical Imaging 2008)*, vol. 6914. SPIE, 2008.
- [75] L. Tang, M. Garvin, K. Lee, W. Alward, Y. Kwon, and M. Abramoff, “Robust Multi-Scale Stereo Matching from Fundus Images with Radiometric Differences,” *IEEE Trans. Pattern Anal. Mach. Intell.*, no. 99, p. 1, 2011.
- [76] K. A. Vermeer, J. van der Schoot, H. G. Lemij, and J. F. de Boer, “Automated segmentation by pixel classification of retinal layers in ophthalmic OCT images,” *Biomed. Opt. Express*, vol. 2, no. 6, pp. 1743–56, Jun. 2011.
- [77] R. M. Haralick, K. Shanmugam, and I. Dinstein, “Textural features for image classification,” *IEEE Trans. Syst., Man, Cybern.*, vol. 3, no. 6, pp. 610–621, 1973.
- [78] M. M. Galloway, “Texture analysis using gray level run lengths,” *Computer Graphics and Image Processing*, vol. 4, no. 2, pp. 172–179, 1975.
- [79] P. Viola, O. M. Way, and M. J. Jones, “Robust Real-Time Face Detection,” *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.
- [80] S. E. Grigorescu, N. Petkov, and P. Kruizinga, “Comparison of texture features based on Gabor filters,” *IEEE Transactions on Image Processing*, vol. 11, no. 10, pp. 1160–7, Jan. 2002. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18249688>
- [81] B. J. Antony, M. D. Abramoff, M. Sonka, Y. H. Kwon, and M. K. Garvin, “Incorporation of texture-based features in optimal graph-theoretic approach with application to the 3-D Segmentation of intraretinal surfaces in SD-OCT volumes,” in *Proc. SPIE 2012*, D. R. Haynor and S. Ourselin, Eds., vol. 8314, San Diego, CA, 2012, p. 83141G.

- [82] W. T. Freeman and E. H. Adelson, “The design and use of steerable filters,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 9, pp. 891–906, 1991.
- [83] P. Pudil, J. Novovičová, and J. Kittler, “Floating Search Methods In Feature Selection with Nonmonotonic Criterion Functions,” *Pattern Recognition Letters*, vol. 15, no. 11, pp. 1119–1125, Nov. 1994.
- [84] D. C. Lozano and M. D. Twa, “Quantitative evaluation of factors influencing the repeatability of SD-OCT thickness measurements in the rat,” *Invest. Ophthalmol. Vis. Sci.*, vol. 53, no. 13, pp. 8378–85, Dec. 2012.
- [85] Y. Muraoka, H. O. Ikeda, N. Nakano, M. Hangai, Y. Toda, K. Okamoto-Furuta, H. Kohda, M. Kondo, H. Terasaki, A. Kakizuka, and N. Yoshimura, “Real-time imaging of rabbit retina with retinal degeneration by using spectral-domain optical coherence tomography.” *PLoS One*, vol. 7, no. 4, p. e36135, Jan. 2012.
- [86] T. J. Bailey, D. H. Davis, J. E. Vance, and D. R. Hyde, “Spectral-domain optical coherence tomography as a noninvasive method to assess damaged and regenerating adult zebrafish retinas.” *Invest. Ophthalmol. Vis. Sci.*, vol. 53, no. 6, pp. 3126–38, May 2012.
- [87] M. E. Pennesi, K. V. Michaels, S. S. Magee, A. Maricle, S. P. Davin, A. K. Garg, M. J. Gale, D. C. Tu, Y. Wen, L. R. Erker, and P. J. Francis, “Long-term characterization of retinal degeneration in rd1 and rd10 mice using spectral domain optical coherence tomography.” *Invest. Ophthalmol. Vis. Sci.*, vol. 53, no. 8, pp. 4644–56, Jul. 2012.
- [88] G. Huber, S. C. Beck, C. Grimm, A. Sahaboglu-Tekgoz, F. Paquet-Durand, A. Wenzel, P. Humphries, T. M. Redmond, M. W. Seeliger, and M. D. Fischer, “Spectral domain optical coherence tomography in mouse models of retinal degeneration.” *Invest. Ophthalmol. Vis. Sci.*, vol. 50, no. 12, pp. 5888–95, Dec. 2009.
- [89] M. L. Gabriele, H. Ishikawa, J. S. Schuman, R. A. Bilonick, J. Kim, L. Kagemann, and G. Wollstein, “Reproducibility of spectral-domain optical coherence tomography total retinal thickness measurements in mice.” *Invest. Ophthalmol. Vis. Sci.*, vol. 51, no. 12, pp. 6519–23, Dec. 2010.
- [90] W. J. Jeong, M. D. Abramoff, B. J. Antony, C. Jiao, M. K. Garvin, and E. H. Sohn, “Retinal thickness changes in mice with streptozotocin-induced diabetes mellitus quantified using an enhanced Iowa reference algorithm,” *Invest. Ophthalmol. Vis. Sci. E-Abstract*, p. 4883, 2013.

- [91] B. J. Antony, M. D. Abramoff, W. Jeong, E. H. Sohn, and M. K. Garvin, "Segmentation of Multiple Intra-retinal Surfaces in Volumetric SD-OCT Images of Mouse Eyes Using an Improved Iowa Reference Algorithm," *Invest. Ophthalmol. Vis. Sci. E-Abstract*, vol. 54, p. 4892, 2013.
- [92] N. Nakano, H. O. Ikeda, M. Hangai, Y. Muraoka, Y. Toda, A. Kakizuka, and N. Yoshimura, "Longitudinal and simultaneous imaging of retinal ganglion cells and inner retinal layers in a mouse model of glaucoma induced by N-methyl-D-aspartate," *Invest. Ophthalmol. Vis. Sci.*, vol. 52, no. 12, pp. 8754–62, Jan. 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22003119>
- [93] A. S. C. Reis, G. P. Sharpe, H. Yang, M. T. Nicolela, C. F. Burgoyne, and B. C. Chauhan, "Optic disc margin anatomy in patients with glaucoma and normal controls with spectral domain optical coherence tomography." *Ophthalmology*, vol. 119, no. 4, pp. 738–47, Apr. 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22222150>
- [94] K. L. Boyer, A. Herzog, and C. Roberts, "Automatic recovery of the optic nerve-head geometry in optical coherence tomography," *IEEE Trans. Med. Imag.*, vol. 25, no. 5, pp. 553–570, 2006.
- [95] G. P. Nason and B. W. Silverman, "The Stationary Wavelet Transform and some Statistical Applications," *Wavelets and statistics*, pp. 281–299, 1995.
- [96] D. Marr and E. Hildreth, "Theory of Edge Detection," *Proc. of the Royal Society*, vol. 207, no. 1167, pp. 187–217, 1980.
- [97] C. Harris and M. Stephens, "A Combined Corner and Edge Detector," *Proceedings of the Alvey Vision Conference 1988*, vol. 15, p. 50, 1988. [Online]. Available: <http://www.bmva.org/bmvc/1988/avc-88-023.html>
- [98] S. M. Smith and J. M. Brady, "SUSAN - A New Approach to Low Level Image Processing," *International Journal of Computer Vision*, vol. 23, pp. 45—78, 1995.
- [99] D. Freedman and T. Zhang, "Interactive graph cut based segmentation with shape priors," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, Jun. 2005, pp. 755–762.
- [100] Z. Hu, "Multimodal 3-D segmentation of optic nerve head structures from spectral domain OCT volumes and color fundus photographs," Ph.D. dissertation, The University of Iowa, 2011.

- [101] B. C. Chauhan and C. F. Burgoyne, "From clinical examination of the optic disc to clinical assessment of the optic nerve head: a paradigm change." *American journal of ophthalmology*, vol. 156, no. 2, pp. 218–227.e2, Aug. 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23768651>
- [102] N. W. Nickells, "Ganglion cell death in glaucoma: mice to men," *Veterinary Ophthalmology*, vol. 10, pp. 88–94, 2007.
- [103] X. Zhang, C. J. Bregman, A. S. Raza, G. De Moraes, and D. C. Hood, "Deriving visual field loss based upon OCT of inner retinal thicknesses of the macula," *Biomed. Opt. Express*, vol. 2, no. 6, pp. 1734–1742, Jun. 2011.
- [104] S. D. Grozdanic, H. Kecova, M. M. Harper, W. Nilaweera, M. H. Kuehn, and R. H. Kardon, "Functional and structural changes in a canine model of hereditary primary angle-closure glaucoma." *Investigative ophthalmology & visual science*, vol. 51, no. 1, pp. 255–63, Jan. 2010.