Fall 2009

Automated 3-D segmentation of intraretinal surfaces from optical coherence tomography images centered on the optic nerve head

Bhavna Josephine Antony
University of Iowa

AUTOMATED 3-D SEGMENTATION OF INTRARETINAL LAYERS FROM OPTICAL COHERENCE TOMOGRAPHY IMAGES CENTERED ON THE OPTIC NERVE HEAD

by

Bhavna Josephine Antony

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

December 2009

Thesis Supervisor: Assistant Professor Mona K. Garvin

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

MASTER'S THESIS

This is to certify that the Master's thesis of

Bhavna Josephine Antony

has been approved by the Examining Committee for the thesis requirement for the Master of Science degree in Electrical and Computer Engineering at the December 2009 graduation.

Thesis Committee: _____
Mona K. Garvin, Thesis Supervisor

_____
Michael D. Abràmoff

_____
Milan Sonka

To Kelly.

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 for all his advice and encouragement during my time here at the University of Iowa, and Dr. Michael Abramoff for his invaluable inputs.

I an extremely grateful to Dr. Priya Gupta, Pavlina Sonkova and Dr. Mona Garvin for the time and the efforts they spent on manually tracing the surfaces on the OCT images. The creation of the reference set is an extremely time consuming process and I am very grateful to all of them.

I extend special thanks to my student colleagues Zhihong Hu, Kit Lee and Li Tang for all their help and support.

I also thank my parents and elder sister for being such endless sources of encouragement, and for giving me the opportunity to come here to the University of Iowa for my graduate studies. I could not have achieved any of this without their unconditional support.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

CHAPTER 1
INTRODUCTION

Since its introduction in 1991 [1], optical coherence tomography has found wide application as an important non-invasive imaging modality, notably in ophthalmology where it is used to image the retina and the optic nerve head. These high resolution images have generated a better understanding of the structural consequences of ocular diseases and as such, has found increasing application in the diagnosis and management of numerous diseases. The early generation scanners were time-domain systems that acquired up to six circular/radial scans in 2-D. The latest generation of scanners, however, are Fourier-domain systems (which became commercially available in 2007) that possess several advantages over the older time-domain scanners. Not only do the Fourier-domain OCT images (FD-OCT) provide true volumetric data, the scans also boast higher resolutions (5μm vs. 8-10μm in time-domain OCT) making the FD-OCT images considerably larger than their time-domain counterparts. For instance, a time-domain dataset acquired using a 6-radial scan protocol (obtained on the Stratus OCT-3, Carl Zeiss Meditec, Inc., Dublin) has a size of $6 \times 128 \times 1024$ pixels while many FD-OCT images (obtained on the Cirrus, Carl Zeiss Meditec) have dimensions of $200 \times 200 \times 1024$ pixels.

The need for automated methods grows with the increasing availability of data to ophthalmologists, but at present, on optic nerve head (ONH) scans the commercial scanners only segment the retinal nerve fiber layer (RNFL), potentially disregarding valuable indicators from the other layers. It is also likely that the method used for this segmentation is a 2-D approach that does not use any of the 3-D contextual information available in the new volumetric data. This is also the case with most of the retinal layer segmentation methods that have been proposed thus far. However, the advantages of segmenting medical volumetric images in 3-D is balanced by the complexity of the problem and the difficulty faced in reaching an optimal solution,

which grows with the size of the dataset. In this regard, the method described by Li *et. al* [2] is of significance, as it segments multiple surfaces simultaneously from an n-D volume, assures optimality and is a low polynomial-time algorithm. This graph-based method transforms the segmentation problem into an optimization problem that aims to find a feasible set of surfaces with the minimum cost.

The advantages of such a segmentation method has been amply established by Garvin [3] and Lee [4], who used such an approach for the segmentation of retinal layers at the macula. Garvin [3] was the first to propose the use of a graph-based approach for the simultaneous segmentation of the retinal surfaces, and also proposed additions to the original formulation of the problem in the form of varying feasibility constraints that are learned from a training set, and the incorporation of "true" regional information in the cost function. The method described by Lee [4] on the other hand, did not incorporate any of these changes, but uses a multi-resolution approach in which the surfaces are segmented in a number of graph-searches instead of being detected simultaneously. The method's key advantage is its small run-time of a few minutes.

In this thesis, we present a method for the segmentation of retinal layers from OCT scans centered on the ONH, a region that has thus far, remained largely unexplored. To the best of our knowledge, there are no known validated techniques that currently segment the layers at the ONH region. The method described herein combines key aspects of the segmentation techniques proposed by Garvin *et. al* [3] and Lee [4], and uses a multi-resolution approach, while also incorporating learned varying feasibility constraints and regional information. We have also proposed and validated a new technique for the flattening of the OCT datasets, and provided a comparison to the approach adopted by Garvin *et. al* [3]. The challenges posed by the neural canal and the large blood vessels characteristic to this region have also been addressed. In addition to the validation on normal and as well as diseased (glaucomatous) eyes, we

also present some measurements that could be clinically relevant.

CHAPTER 2
BACKGROUND

2.1 The Retina

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.

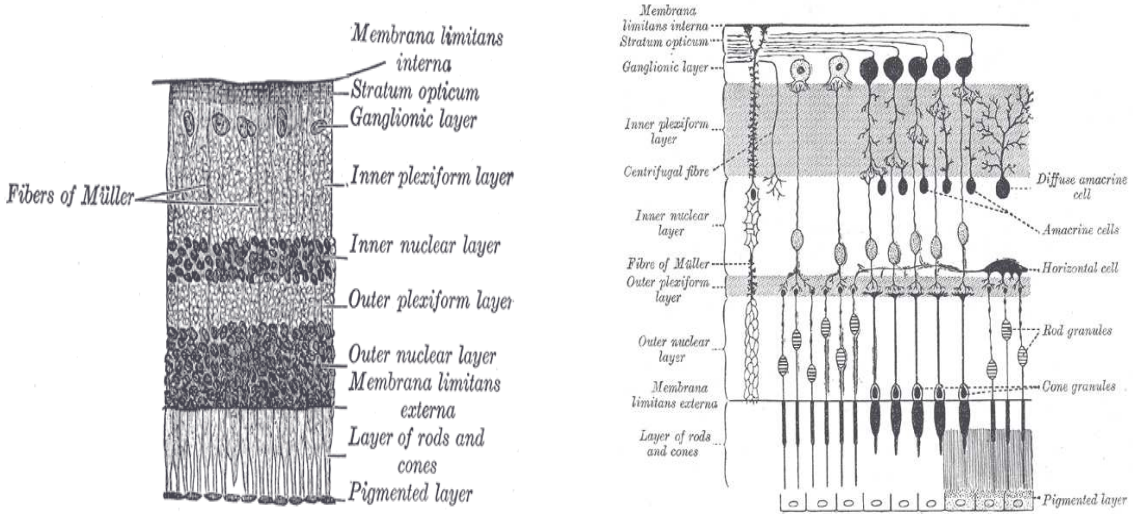

(a) (b)

Figure 2.1: Lithograph plates showing the layers of the retina. These images from Gray's Anatomy (originally published in 1918 and has since lapsed into public domain) show (a) a cross-sectional view of the ten retinal ten layers, (b) the various cells found in the retina.

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 Cells 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.
- External Limiting Membrane (ELM) separates the inner-outer segments of the photoreceptors (both rods and cones) from their cell bodies.
- Photoreceptor Layer contains the segments (light reactive parts) of the photoreceptors (both rods and cones) and is roughly divided into two layers, the inner segments (IS) and the outer segments (OS).
- Retinal Pigment Epithelium (RPE) is a single layer of cells between the retina and the choroid.

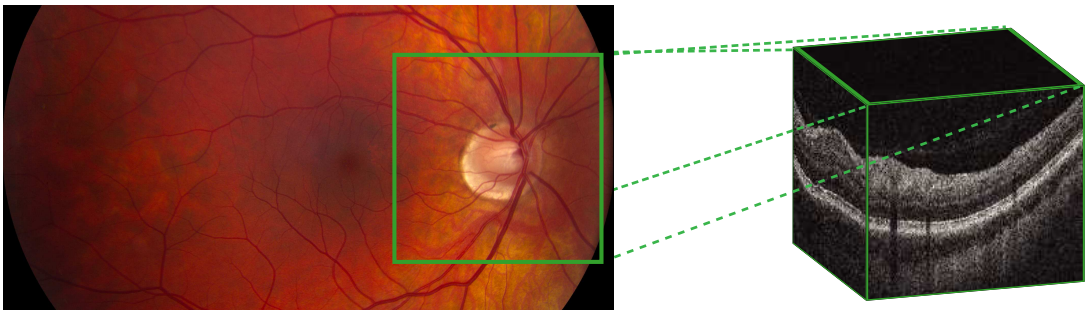

(a)

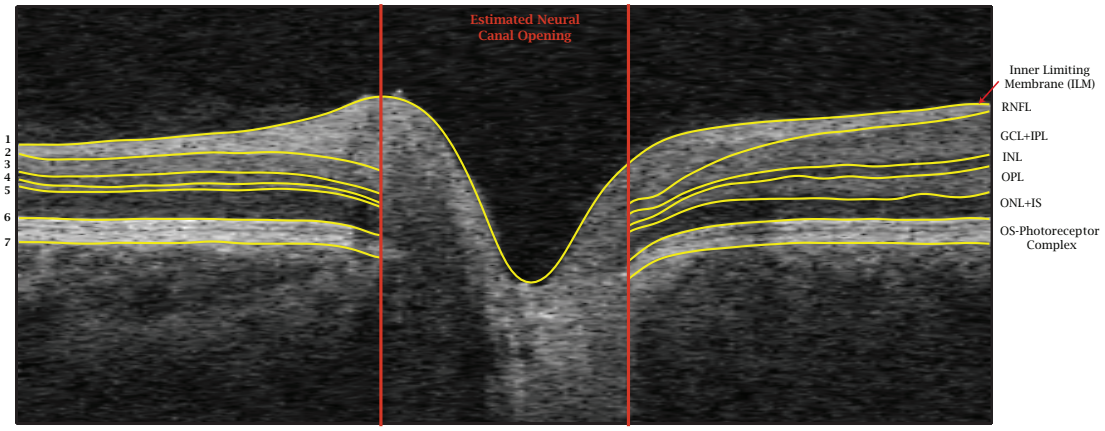

(b)

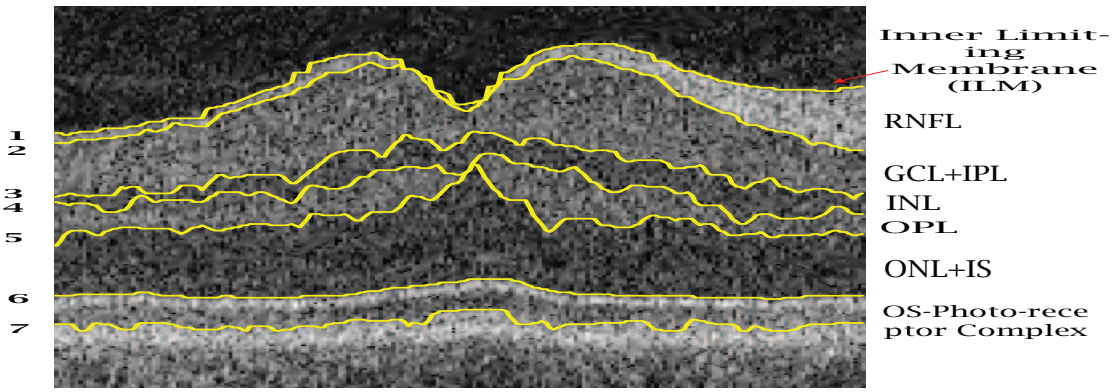

(c)

Figure 2.2: OCT images from the optic nerve head. (a) Fundus Photograph showing the location of OCT scans with respect to a retinal fundus photograph. (b) A central xz-slice from the 3-D OCT dataset showing the various layers of the retina as well as the neural canal opening. (c) Shows the surfaces on a central xz-slices from a macular OCT dataset.

In the OCT images, the surface between the GCL and IPL layers is not consistently seen and is therefore, not segmented. The same is true of the ELM as well. Thus, the 6 layers that are segmented include the RNFL, the combined GCL and IPL, INL, OPL, ONL (which includes the ELM) and the OS-Photoreceptor complex (Fig. 2.2). The optic nerve head (Fig. 2.2(a)) 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 show a number of large blood vessels as this is the location from which the vessels that supply the ocular tissue enter and leave the eye.

2.2 Clinical Motivation

In this work, we focus on the possible clinical applications of peripapillary layer segmentation for the diagnosis and management of glaucoma. This disease is the most common optic nerve disorder, and affects 1-2% of the U.S. population [5,6] and is the second leading cause of blindness worldwide. Glaucoma is often characterized by an increase in intraocular pressure (IOP) and in the past, was often defined by this symptom. But recent studies have showed 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. Visual field tests are good indicators of visual acuity, but the reproducibility is poor in patients that have suffered moderate to severe visual field loss. Thus, reliable methods that can diagnose and monitor the progression of the disease are extremely important. Furthermore, structurally the disease is known to affect the nerve fiber layer which is discernible in

OCT images. The thinning of the RNFL is an important marker used by ophthalmologists in the diagnosis of the disease. It has also been used to differentiate between glaucomatous and normal eyes [7–10]. Despite the fact that this is a known marker of the disease, the commercial scanners only provide quantitative measurements along a circular path around the optic disc. Further information solely consist of qualitative indicators of possible thinning. The presence of further indicators within the quadrants [11] is also known, but again, is not available to clinicians at present. Measurements that could reveal vital clues include volumetric measurements in the whole scan, within selected sectors/quadrants and radial distances. In this work, we present some of these measurements and show the possible clinical applications of a tool that provides complete layer segmentations from OCT datasets.

2.3 Previous Work

Most of the previously reported intraretinal layer segmentation approaches deal solely with the macular region and were designed to deal with the smaller time-domain datasets. The commerically available scanners do provide some information about the RNFL in the form of thickness at a particular radial distance from the center of the volume. The scanners also indicate regions that show considerable thinning, but provide no actual quantitative measurements in this regions or volumetric measurements of the whole layer. It is also likely that the method used, like the majority of the proposed segmentation methods, is a 2-D approach. [9, 12–16] proposed retinal layer segmentation techniques that rely heavily on edge information in each A-scan line and use 2-D contextual information to verify the detected edges. For instance, the methods [9, 12–15] follow four main steps. 1) The datatsets are filtered, 2) edges are detected in each A-scan line, 3) the large gradient that marks the location of surface 6 (Fig.2.2) 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. [16] proposed the mapping of grayscale intensities as an additional

feature, but the method's ability to deal with diseased datasets where the images sometimes show poor contrast was not adequately demonstrated.

More recently, [3,4] have presented graph-based 3-D approaches for the segmentation of intraretinal layers at the macula. Garvin *et. al* [3] was the first to propose the use of a 3-D graph-search to segment multiple surfaces from retinal OCT scans. This approach incorporated two key additions to the original 3-D graph-based segmentation method [2]. 1) The feasibility constraints used were learned from a training set and varied as a function of (x, y) location, and 2) the cost function incorporated "true" regional information. Lee [4] also proposed a segmentation method that aimed to deal with the large memory and run-time requirements of the previous method, and differs in two aspects. 1) The cost function is derived from edge gradients and does not use any form of regional cost terms, and 2) the segmentation method itself is a multi-resolution approach that does not segment all the retinal surfaces simultaneously, but uses multiple graph-searches to obtain the final result. These two methods used thin-plate splines to flatten the datasets, but Garvin *et. al* [3] used a 3-D thin-plate spline with a smoothing regularization term, that flattened the dataset with respect to a smoothed surface, while Lee [4] used a 2-D spline with a very small regularization term that flattened each B-scan separately.

CHAPTER 3
3-D GRAPH-BASED SURFACE DETECTION

In this chapter, we review graph search concepts as described in [2, 17] and the extensions described in [18, 19].

3.1 Graph-Based Surface Segmentation

The graph-search approach for the segmentation of images, as described in [2,17], transforms the segmentation problem into an optimization problem, where the aim is to find a set of surfaces that form a minimal cost *closed set*. (A *closed set* in a graph is defined as a subset of vertices such that no directed edges of the graph leave the set.) The transformation of the problem begins with the construction of a vertex-weighted graph which uses two main descriptors: constraints that describe the feasiblity of the surfaces and the cost of the surfaces, which are further described below. The optimal solution is then computed on this graph to obtain the set of surfaces that satisfy the predefined constraints.

3.1.1 Feasibility Constraints

Consider a volumetric image $\mathcal{I}(x, y, z)$ of size $X \times Y \times Z$, where each surface $\mathcal{S}$ is defined by a function $\mathcal{S}(x, y)$, where x $\in \mathbf{x} = \{0...$ X-1$\}$, y $\in \mathbf{y} = \{0...$ Y-1$\}$ and $\mathcal{S}$(x,y) $\in \mathbf{z} = \{0...$ Z-1$\}$. Each (x, y) location is associated with a z-column of voxels, and therefore the surface $\mathcal{S}$ intersects each column at a single pixel location.

The feasibility of each surface can be expressed using two parameters, namely the smoothness constraints and the surface interaction constraints. [2, 17] define the smoothness constraints Δ_x and Δ_y, as 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

$$|z_1 - z_2| \leq \Delta_x \ . \tag{3.1}$$

Δ_y is expressed in a similar manner for adjacent voxels in the y-direction. An extension to this is described in [18], where the smoothness constraints are allowed to vary with respect to the (x, y) location of the voxel. The smoothness constraint in the x-direction when moving from location (x_1, y_1) to (x_2, y_2) is now defined as follows:

$$-\Delta^u_{(x_1,y_1),(x_2,y_2)} \leq f(x_1, y_1) - f(x_2, y_2) \leq \Delta^l_{(x_1,y_1),(x_2,y_2)} \ , \tag{3.2}$$

where, $\Delta^u_{(x_1,y_1),(x_2,y_2)}$ and $\Delta^l_{(x_1,y_1),(x_2,y_2)}$ represent the smallest and largest allowed change in z-value, respectively.

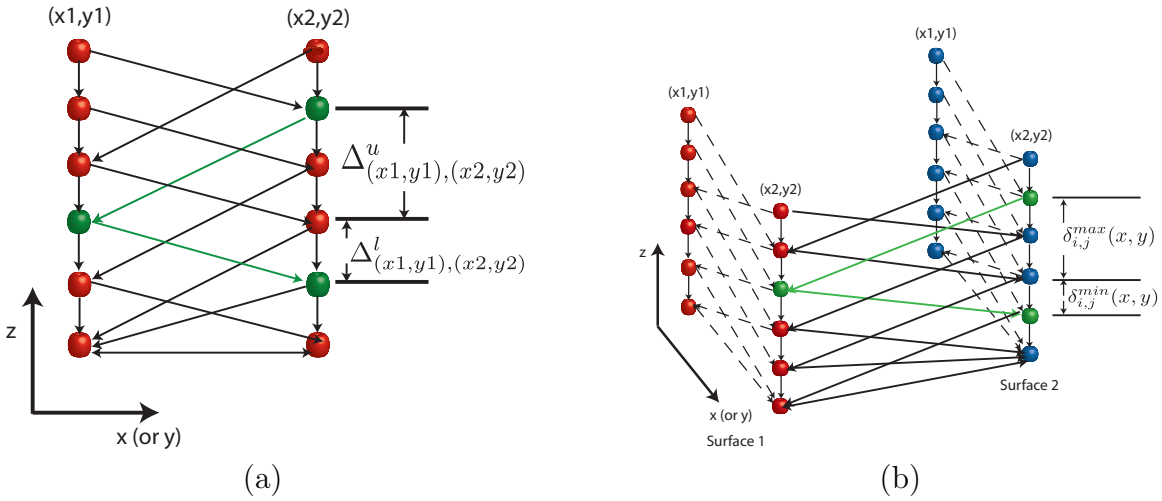

Figure 3.1: Varying feasibility constraints. (a) Shows the varying smoothness constraints. (b) Shows the varying surface-interaction constraints

The second feasibility constraint, the surface-interaction constraint, defines 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 be known that surface i is above surface j and that the minimum and maximum allowed distances between them is $\delta^{min}_{i,j}$ and $\delta^{max}_{i,j}$, respectively. The surface-interaction constraints, like the smoothness

constraints, can be defined as constant values for pairs of adjacent surfaces [2,17] or they can be allowed to vary as a function of the (x, y) location, [3, 18] and can be expressed as follows:

$$\delta_{i,j}^{min}(x,y) \leq \mathcal{S}_j(x,y) - \mathcal{S}_i(x,y) \leq \delta_{i,j}^{max}(x,y) \ , \tag{3.3}$$

where, $\delta_{i,j}^{min}(x,y)$ and $\delta_{i,j}^{max}(x,y)$ represent the minimum and maximum allowed distance between surfaces i and j at location (x, y).

3.1.2 Cost of a Feasible Set of Surfaces

The cost function is an important aspect of the problem description and can be designed using a combination of on-surface and region-based costs. In the original formulation [2], the regional costs were derived from a small region around each voxel. [3, 19] 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 $n + 1$ region based cost terms, as shown in Fig. 3.2. The on-surface cost terms reflect the unlikelihood that the voxel lies on surface i, $i \in \mathbf{i} = \{1...n\}$, while the region cost terms reflect the unlikelihood that the voxel belongs in region i, $i \in \mathbf{i} = \{0....n\}$. If $C_{\mathcal{S}_i(x,y)}$ represents the total cost of surface $\mathcal{S}_i$, then

$$C_{\mathcal{S}_i} = \sum_{(x,y,z)|z=\mathcal{S}_i(x,y)} c_{\mathcal{S}_i(x,y)} \ , \tag{3.4}$$

where, $c_{\mathcal{S}_i(x,y)}$ represents the on-surface costs terms associated with surface i.

Similarly, the "true" regional cost term C_{R_i} that reflects the total costs of the i^{th} region can be defined as:

$$C_{R_i} = \sum_{(x,y,z)\in R_i} c_{reg_i}(x,y,z) \ , \tag{3.5}$$

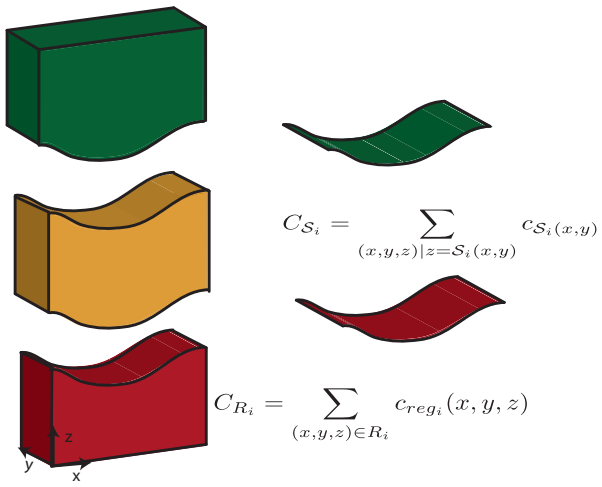

Figure 3.2: Illustration of the regional and on-surface costs terms in a 2 surface segmentation problem. Here, the two surfaces divides the volume into 3 regions.

where, $c_{reg_i}(x, y, z)$ represents the region costs associated with region i. The total cost C_T, associated with the set of n surfaces can now be expressed as a sum of the on-surface and regional cost terms:

$$C_T = \sum_{i=1}^{n} C_{\mathcal{S}_i} + \sum_{i=0}^{n} C_{R_i} \ . \tag{3.6}$$

3.1.3 Graph Construction

A vertex-weighted graph can now be constructed using the cost terms and the feasibility constraints. The weights of the vertices are set so that the cost of each closed set corresponds to the cost of the set of surfaces. [3, 18] defines the weight w_i of each vertex as the sum of two terms, w_{s_i} and w_{r_i}, related to the on-surface and in-region costs, respectively.

The arcs between the nodes are inserted to reflect the pre-defined feasibility constraints. For this, a sub-graph is created for each surface whose structure reflects the smoothness constraints for that particular surface. The sub-graphs are then connected by arcs such that the varying surface-interaction constraints are maintained [2, 3, 18].

Multiple surfaces can now be detected by computing the minimum-cost closed set on this closely related graph, through a minimum-cost $s-t$ cut [2, 17].

CHAPTER 4
SEGMENTATION OF INTRARETINAL SURFACES IN ONH CENTERED OCT SCANS

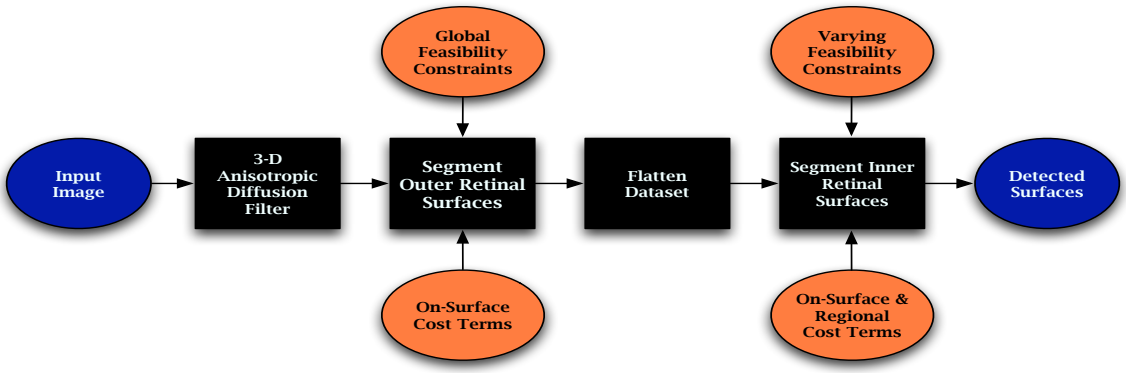

Figure 4.1: Overview of the segmentation method.

In this chapter, we describe the method used to segment 7 retinal surfaces from ONH-centered OCT images. Fig. 4.1 briefly describes the various steps of the segmentation technique. The datasets are first filtered using a 3-D gradient anisotropic diffusion filter as proposed in [20], which helps reduce the speckling characteristic to OCT images. The segmentation can be roughly divided into two stages – the outer surface segmentation and dataset flattening, and the inner surface segmentation – and are described in detail in the following sections.

4.1 Flattening the OCT Volume

The retinal surfaces in the OCT images obtained from the scanner (Fig. 4.2) are far from flat, and in fact distorted in different ways in the B-scan (xz-slices, see Fig. 4.2(a), Fig. 4.4) and C-scans slices (yz-slices, see Fig. 4.2(b)). These artifacts are thought to be the result of a number of factors such as the corneal curvature, motion of the eye and the positioning of the camera. Flattening the dataset can

eliminate these effects, and makes visualization easier by bringing the dataset into a more consistent shape. It also centers the data in a manner that allows for the efficient truncation of the volume to include only the surfaces that are of interest.

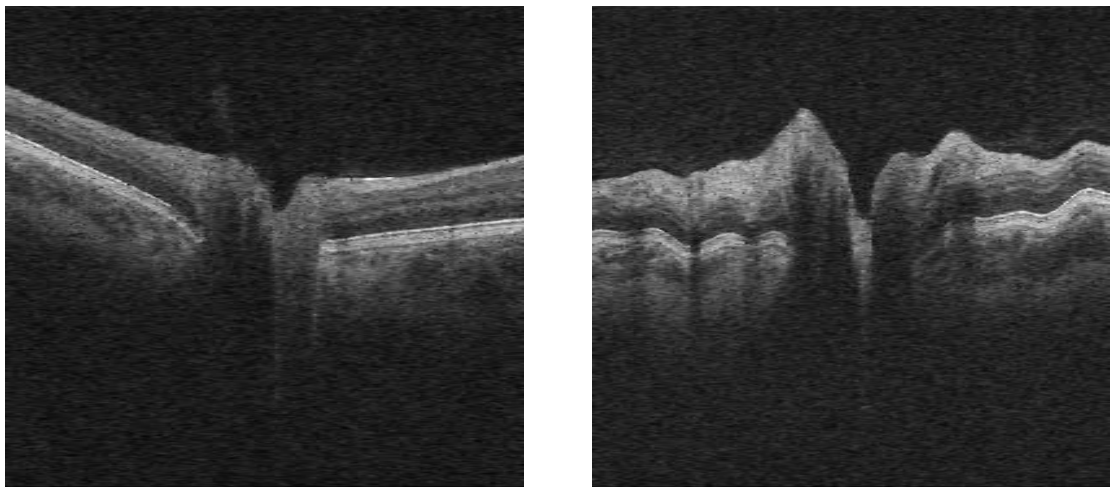

(a) (b)

Figure 4.2: OCT scans before flattening. (a)Central B-scan showing the curvature commonly seen in these slices. (b) Central C-scan showing artifacts characteristic to yz-slices.

The memory requirements of the algorithm are largely dictated by the size of the graph, which in turn depends on the size of the dataset. Thus, flattening and truncating the volume to only include the region of interest offers savings in memory requirements as well as run-time.

As the OCT volumes are flattened with respect to surface 6, the flattening process begins with the detection of the outer surfaces, namely surface 1, 6 and 7 (Fig. 2.2). As depicted in Fig. 4.3, the outer surfaces are found through a multi-resolution approach in 3 steps, at $(1/8)^{th}$, $(1/4)^{th}$ and the full resolution, using the 3-D graph-based approached discussed in the previous chapter. The cost functions and feasibility constraints used in this stage are as described below.

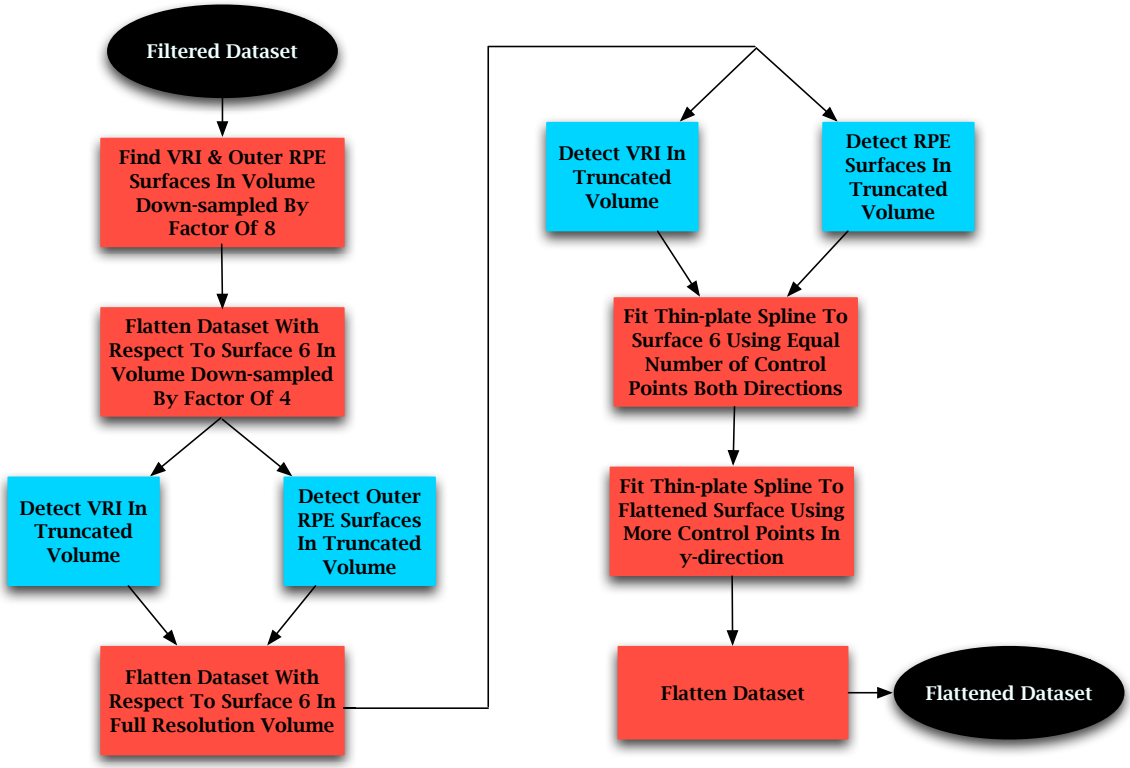

Figure 4.3: Flowchart showing main steps involved in the flattening process.

4.1.1 Construction of the Graph

4.1.1.1 Cost Function

The outer surfaces (surfaces 1, 6 and 7) mark the border between dark and bright regions, therefore, the on-surface costs can be set using gradient information. The gradients are obtained using a combination of vertical and horizontal edge detectors. The horizontal gradient is negligible or absent in most regions of the image, and only provides a response at the neural canal opening (NCO) and vessel shadows. Since the vertical gradients, in these specific areas are extremely small, the horizontal edge information can be used to enhance the overall edge response in the image. Thus, the intensity of the gradient at voxel location (x, y, z) can be expressed as:

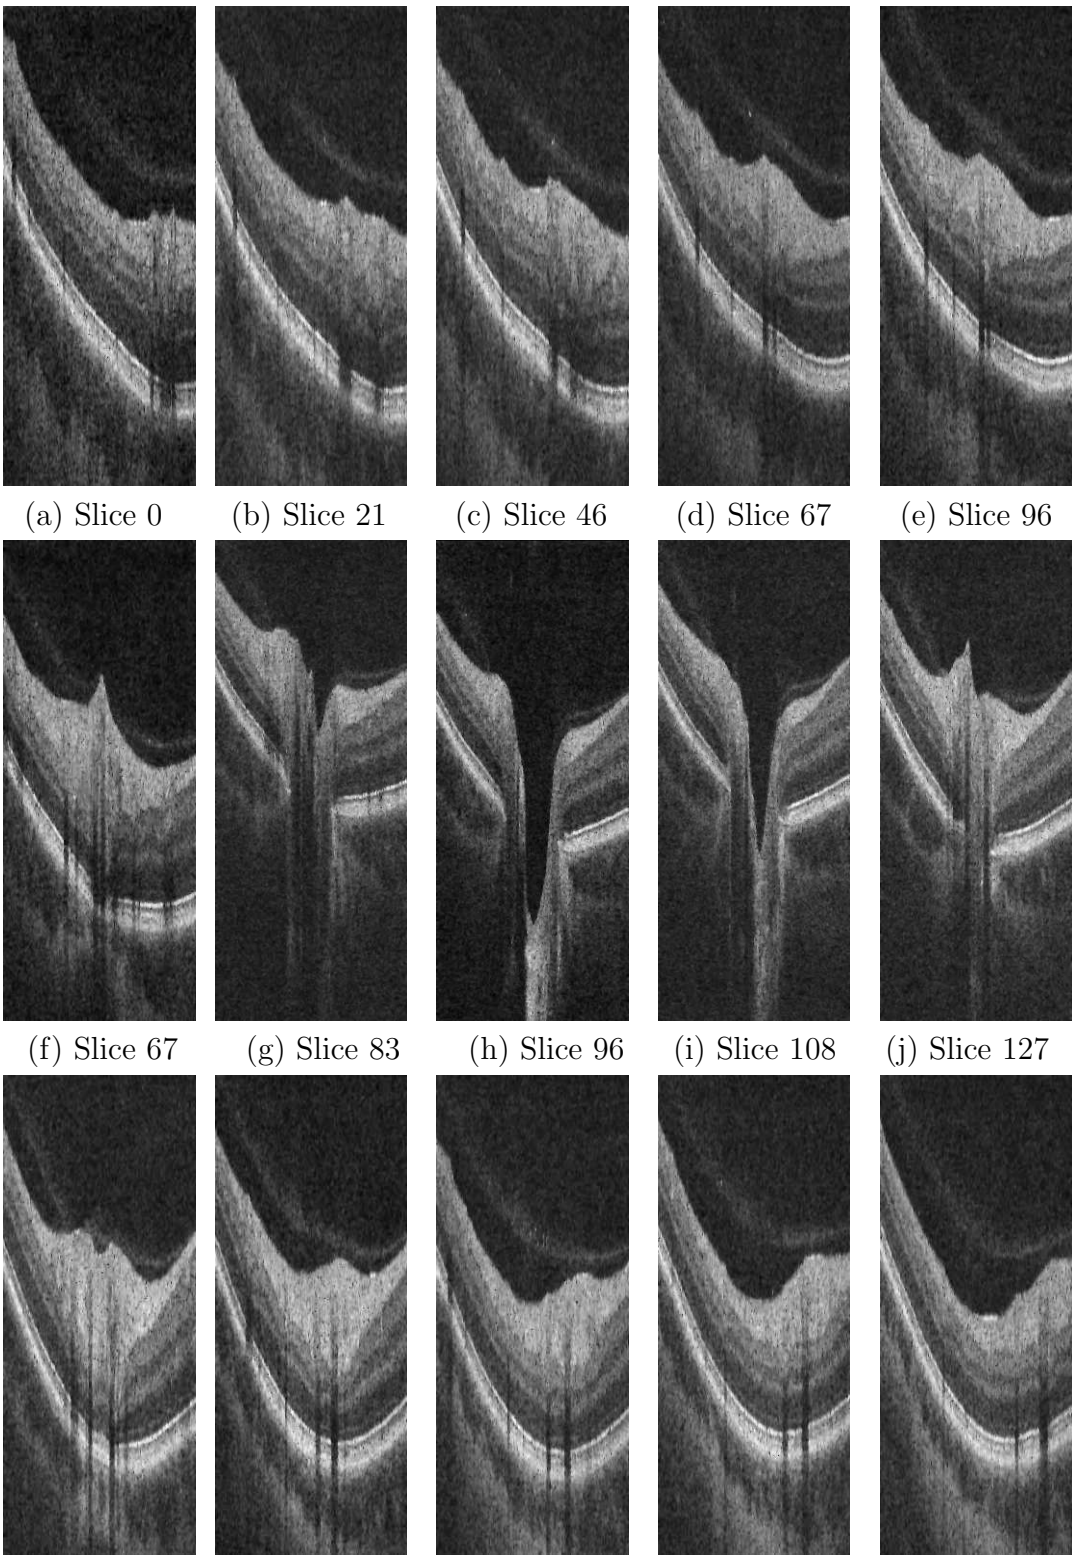

(a) Slice 0 (b) Slice 21 (c) Slice 46 (d) Slice 67 (e) Slice 96

(f) Slice 67 (g) Slice 83 (h) Slice 96 (i) Slice 108 (j) Slice 127

(k) Slice 136 (l) Slice 153 (m) Slice 166 (n) Slice 181 (o) Slice 195

Figure 4.4: Slices from an ONH-centered OCT dataset.

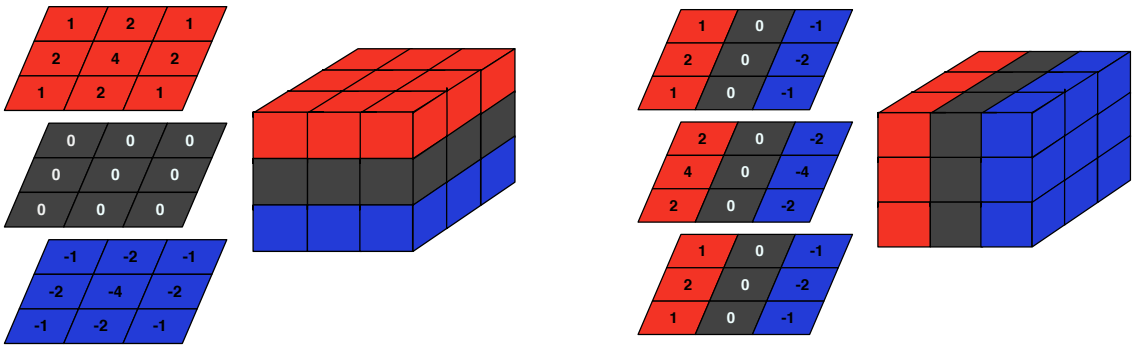

(a) (b)

Figure 4.5: 3-D edge detectors. (a) 3-D vertical edge mask. (b) 3-D horizontal edge mask.

$$G(x, y, z) = \vec{V}(x, y, z) + H(x, y, z) \, , \tag{4.1}$$

where, $\vec{V}(x, y, z)$ represents the directed vertical gradient response and $H(x, y, z)$ represents the absolute horizontal gradient response. Fig. 4.5 shows the 3-D edge detectors used.

Although the gradients are sufficiently strong in the full resolution dataset, in the lower resolutions the images show a considerable amount of noise despite the filtering. But filtering the images further is not an option, as this blurs the edges leading to low gradient responses. Another challenge faced in the lower resolution images, is the large inter-slice variability. Since the dataset has not been flattened yet, the large C-scan artifacts reduce the effectiveness of the 3-D edge detectors. To tackle these two challenges, we use a normalized column-wise cumulative image, where the pixel intensity at (x, y, z) is defined as follows:

$$p(x, y, z) = p(x, y, z) + p(x, y, z - 1), \quad z = \{1, 2.....Z - 1\} \, , \tag{4.2}$$

where, $p(x, y, z)$ is the voxel intensity at the location (x, y, z). The gradients

found on the sum of this cumulative image and the OCT volume are stronger, as the cumulative image incorporates 2-D information and reduces the noise in the image. This cumulative image is used in the detection of the ILM and surface 6. Similarly, the edge magnitudes for surface 7 can be enhanced by using a reversed cumulative image, where:

$$p(x, y, z) = p(x, y, z) + p(x, y, z + 1), \quad z = \{0, 1, 2.....Z - 2\} , \tag{4.3}$$

Thus, if $\mathcal{I}$ represents the image, and $\mathcal{C}$ represents the cumulative image, then

$$\mathcal{I}_c(x, y, z) = \mathcal{I}(x, y, z) + \mathcal{C}(x, y, z), \tag{4.4}$$

$$C_i(x, y, z) = \vec{V}(\mathcal{I}_c(x, y, z)) + H(\mathcal{I}_c(x, y, z)) \tag{4.5}$$

where, C_i is the cost function for the surface $\mathcal{S}_i$.

4.1.1.2 Feasibility Constraints

The feasibility constraints used for the segmentation of the outer surfaces are not derived from a training set, and as such do not vary as a function of the (x, y) location. The photoreceptor layer does not show large variations in thickness and any variability in the thickness between surfaces 1 and 6 is inconsequential, as the ILM and OS-photoreceptor layer complex are found separately in the higher resolutions.

The smoothness constraints are also defined globally, but are adjusted within the neural canal as the surfaces behave differently within this region. The ILM sees large rapid changes as it dips into the neural canal, but surfaces 6 and 7 become indistinct. To compensate for this, the smoothness constraints are relaxed for the ILM and reduced for surfaces 6 and 7 to prevent large (most likely erroneous) variations in the detected surfaces.

4.1.2 Methodology

At the lowest resolution ($(1/8)^{th}$ of the original), the ILM and the two photoreceptor surfaces are segmented simultaneously using a 3-D graph-search. Fig. 4.6(a) shows the three surfaces found on a central xz-slice at this resolution.

The obtained result is used to flatten the dataset in the next resolution (down-sampled by a factor of 4) with respect to surface 6. For this, a thin-plate spline is fit to surface 6 (while avoiding the neural canal by using a circular mask) in order to find a smooth surface at the next (higher) resolution, followed by the translation of all the columns of the dataset such that the spline-fitted surface is reduced to a flat plane. At this resolution, surfaces 6 and 7, and the ILM are found through two separate graph searches. This helps reduces the size of the graphs as the separate graphs only include a small region around the ILM and the OS-photoreceptor complex, where as a simultaneous search for the 3 surfaces would have included a much larger section of the image that contained all the layers of the retina. Separating the two graph-searches also allows us to multi-thread the process, which exploits the multiple cores of the processor and can reduce the processing time by up to a factor of 2 in dual-core processor systems.

In the third and final stage of the segmentation, the result obtained in the previous stage is once more used to localize the graph search in the original resolution and the ILM and surfaces 6 and 7 are again segmented via two graph searches.

Once the outer surfaces have been segmented, the dataset can be flattened by the vertical re-alignment of the columns of the dataset with respect to a plane, which is determined by fitting a thin-plate spline to surface 6. The two main challenges that are faced while flattening are:

- The neural canal, as the surfaces become indiscernible within this region.
- The two distinct distortions seen in the B-scans and the C-scans. The B-scans

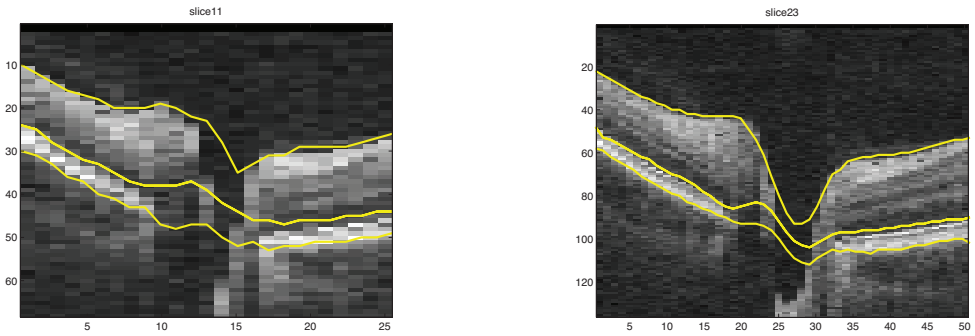

(a) (b)

Figure 4.6: Outer surfaces segmented in lower resolutions. (a) Central B-scan showing the intermediate result at the lowest resolution ($(1/8)^{th}$ original image). (b) Segmentation result at the next resolution ($(1/4)^{th}$ original image).

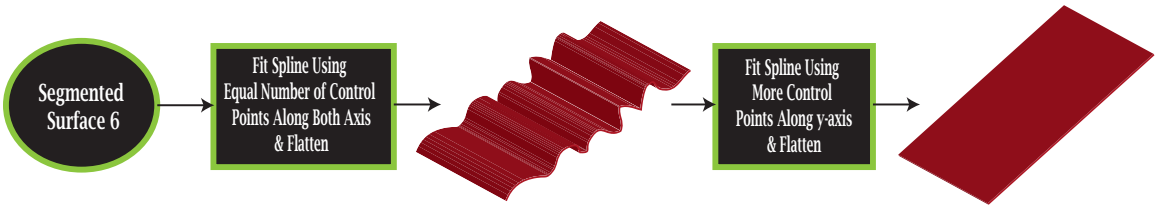

Figure 4.7: Overview of method to determine flattening plane. The flattening plane is determined by fitting a spline to surface 6 twice, to eliminate the two distinct artifacts seen in these images.

show a smooth curvature, while the layers in C-scans sometimes show large ripples.

Control points from within the neural canal can be avoided by using a circular mask to approximate the region. Thus, the flattening plane is not influenced by the unreliability of the detected surfaces within this region. To compensate for the two different artifacts seen in the dataset, the flattening is done in two steps:

- A thin-plate spline is fit to surface 6 using equal number of control points in both of the axial directions, and a smoothing regularization term of 0.1. The

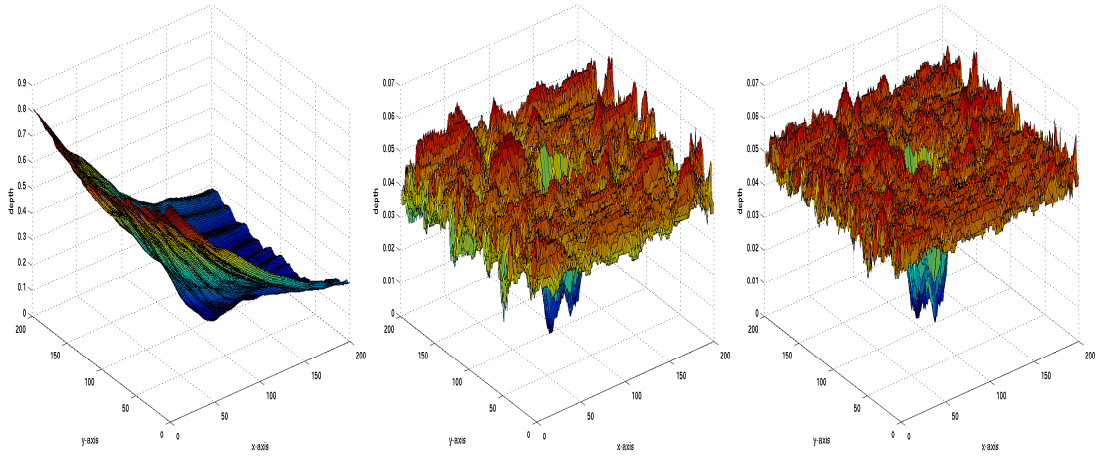

(a) (b) (c)

Figure 4.8: Construction of the flattening plane. (a) Surface 6 as segmented from the original unflattened dataset. (b) Surface 6 after the first stage of flattening with respect to the plane found after the first spline-fit. (c) The final flattened surface obtained after the second spline-fit.

dataset is flattened with respect to this flattening plane.

- A thin-plate spline is fit to the flattened surface using a larger number of control points in the y-direction than the x-direction, and the dataset is flattened once more with respect to the new flattening plane. The regularization term used in this step is a little smaller (0.07) to better approximate the rapid variations seen along the y-direction.

As Fig. 4.7 shows, the thin-plate spline found in the first stage is sufficient to approximate the overall curvature of the dataset, but is not capable of handling the high frequency rippling seen in the C-scans, thus, requiring a second spline fit. Fig. 4.8 shows surface 6 from an example dataset at the various stages of flattening. The first spline-fit, fits a gross smooth surface to the segmented surface 6. As depicted in Fig. 4.8(b), this smooth approximation of surface 6 is not sufficient to completely remove the artifacts seen in the y-direction. But as shown in Fig. 4.8(c), the second-spline

fit which uses a larger number of control points as well as a smaller regularization factor is capable of doing so.

4.2 Segmentation of Inner Retinal Surfaces

The four inner surfaces, like the outer surfaces, (Fig. 4.9) are detected using a multi-resolution approach, but the graph constructed shows two key changes:

- The overall cost function incorporates regional cost terms in addition to on-surface cost terms.
- The feasibility constraints vary as a function of the (x, y) location, and are determined *a priori* from a training set.

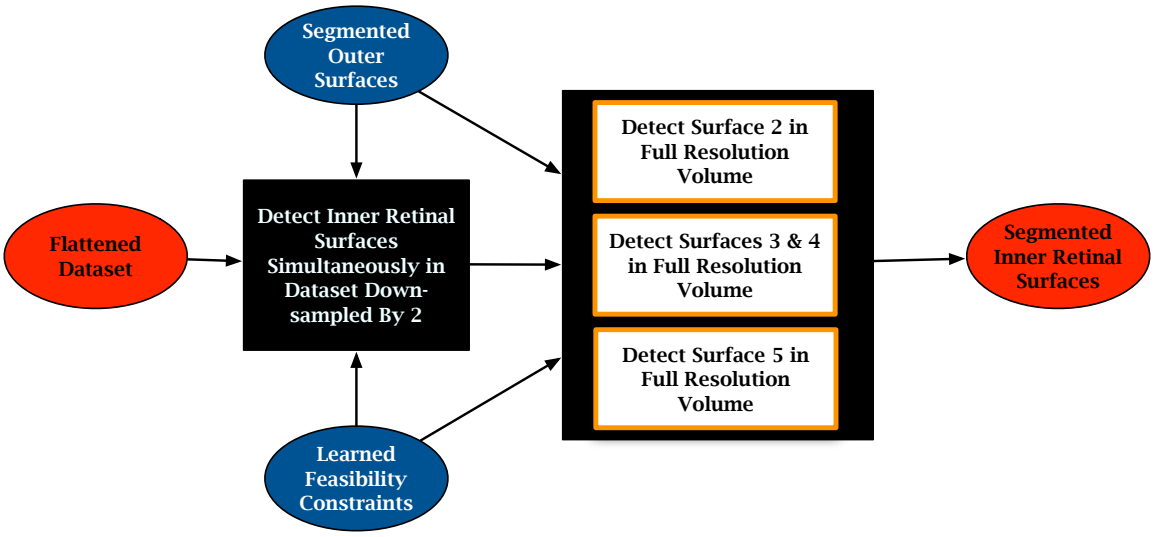

Figure 4.9: Flowchart showing the main steps for the segmentation of the inner retinal surfaces.

4.2.1 Construction of the Graph

4.2.1.1 Cost Function

The cost functions, as described in section 3.1.2, can be defined as combination of on-surface cost terms and regional cost terms. The on-surface cost terms are computed using the 3-D edge detectors as described in section 3.1.1.1. Note that for the inner retinal surfaces, the cumulative image cannot be used as it tends to blur the gradients (which are smaller in comparision to those seen in the outer layers). Thus, the on-surface cost terms are computed as the sum of the directed vertical and the absolute horizontal gradients.

The regional cost terms reflect the likelihood that a pixel belongs to a certain region. For this, we use Gaussian functions as described in [19], to map the normalized intensity values to a likelihood of belong to a particular class. On inspection of the images, we see that the layers can be roughly classified into dark, bright and medium classes. Thus, the in-region cost terms can be derived from these fuzzy memberships. For instance, let $f_d(x)$ represent the probability that a pixel belongs to the dark class, then

$$f_d(x) = \begin{cases} 1 & \text{for x} \leq \ \Delta_d \\ e^{(x-\Delta_d)^2/2\sigma^2} & \text{for x} > \Delta_d \end{cases} \tag{4.6}$$

where, Δ_d is the threshold for the dark class. Similarly, functions are defined for the medium and bright classes.

$$f_m(x) = \begin{cases} e^{(x-(c_m-\Delta_m))^2/2\sigma^2} & \text{for x} \leq c_m - \Delta_m \\ 1 & \text{for } c_m - \Delta_m < \text{x} < c_m + \Delta_m \\ e^{(x-(c_m+\Delta_m))^2/2\sigma^2} & \text{for x} \geq c_m + \Delta_m \end{cases} \tag{4.7}$$

$$f_b(x) = \begin{cases} e^{(x-(1-\Delta b))/2\sigma^2} & \text{for x} < 1 - \Delta_b \\ 1 & \text{for x} > 1 - \Delta_b \ . \end{cases} \tag{4.8}$$

To allow for variability between images, the values of Δ_d, Δ_m, Δ_b are derived from the images. At this stage, the location of the outer surfaces is known, and therefore, can be in conjunction with the learned expected thicknesses to assess average intensities within the three region. Since the ONL is a dark layer, Δ_d can be computed as the mean intensity of voxels between surface 6 and the minimum learned distance between surfaces 5 and 6. Similarly, Δ_b is set to the mean value between surface 1 and the learned minimum distance between surfaces 1 and 2, and Δ_m is set to the average intensity between the minimum distance from surface 1 to surface 3, and maximum distance between surface 1 and surface 3. Fig. 4.10 shows the limits used to compute these values. c_m and σ are set to appropriate small constants.

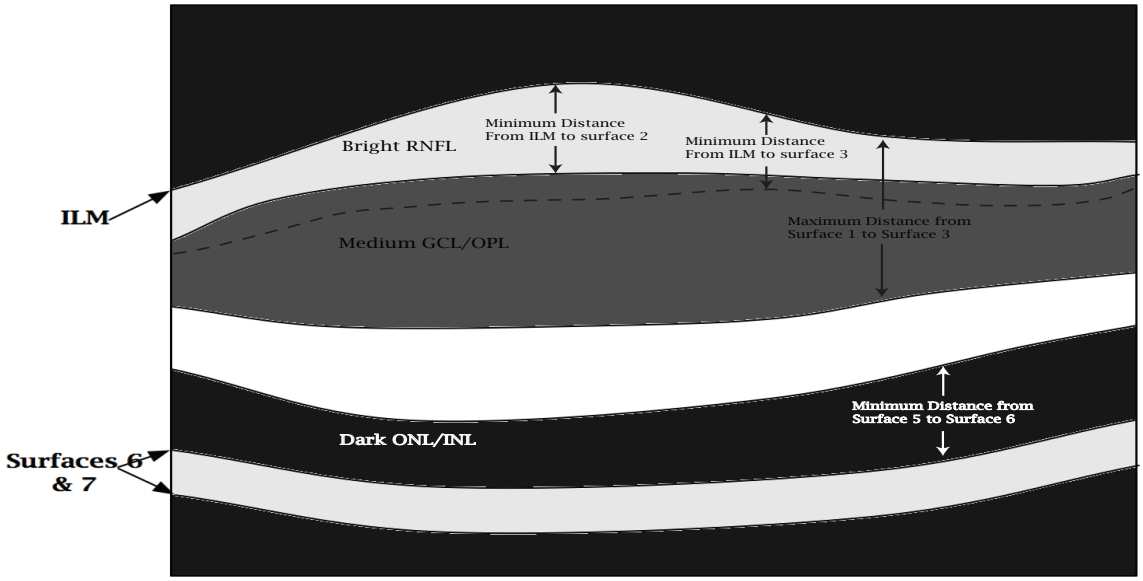

Figure 4.10: The bounds used for the estimation of regional parameters. The region used for the estimation of Δ_d, Δ_m and Δ_b is fixed using the outer surface segmentation result and the learned thickness constraints.

Thus, the overall cost function can now be expressed as the weighted sum of the regional and on-surface cost terms:

$$C_T = \alpha \sum_{i=1}^{n} C_{\mathcal{S}_i} + (1-\alpha) \sum_{i=0}^{n} C_{R_i} \ , \tag{4.9}$$

where, α_i determines the ratio between the regional and on-surface cost terms for surface i. The α_i values used and the method of their determination is further discussed in Section 4.2.3.

4.2.1.2 Feasibility Constraints

As described in Section 3.1, the constraints that describe the feasible set of surfaces can be allowed to vary as a function of the (x, y) location and were learned from a manually traced training set (see Section 5.1). For a four-neighbor system, the varying smoothness constraints are calculated in the x and y directions, each of which are expressed using a minimum $\Delta^l_{(x_1,y_1),(x_2,y_2)}$, and a maximum $\Delta^u_{(x_1,y_1),(x_2,y_2)}$, allowed change in the z-value when moving from (x_1, y_1) to (x_2, y_2). These limits are calculated using the mean and standard deviation of the changes in z-values noted in the training set. First the mean μ_s and standard deviation σ_s are calculated, then the two limits are set as follows:

$$\Delta^l_{(x_1,y_1),(x_2,y_2)} = \mu_s - 3.2 * \sigma_s, \tag{4.10}$$

$$\Delta^u_{(x_1,y_1),(x_2,y_2)} = \mu_s + 3.2 * \sigma_s \ . \tag{4.11}$$

The value of 3.2 is chosen, so that these limits are large enough to account for at least 99% of expected changes in z-values.

Similarly, the varying thickness constraints are determined for each (x, y) location, using the mean and standard deviation of the thicknesses encountered in the training set.

$$\delta_{i,j}^{min}(x,y) = \mu_t - 3.2 * \sigma_t, \tag{4.12}$$

$$\delta_{i,j}^{max}(x,y) = \mu_t + 3.2 * \sigma_t \ . \tag{4.13}$$

Care is taken to make sure that $\delta_{i,j}^{min}$ does not assume values less than 0.

4.2.2 Methodology

The segmentation of the inner surfaces begins with the simultaneous segmentation of the four inner surfaces in the lower resolution (1/2 of the original) flattened volume. The graph is constructed as described in Sec. 3.2.1.1-2, and the four inner surfaces are found simultaneously.

In the full resolution, the four surfaces are segmented via three separate graph-searches (which can be run simultaneous by multi-threading the processes). Surface 5 is detected in a graph constructed from a small region around the surface detected in the lower resolution. Similarly, surface 2 is segmented as a single surface, while surfaces 3 and 4 are segmented together due to their close proximity to one another.

The lower resolution result not only provides information about the location of the surface, it can also be used to provide much better estimates for the regional cost term parameters, Δ_d, Δ_m and Δ_b. The parameters are now computed from regions that are specified by the approximate locations of the surfaces found in the lower resolution. For instance, Δ_d is now computed from the region demarcated by surfaces 6 and the lower resolution approximate location of surface 5. Similarly, Δ_m is set to the mean intensity between the previously detected surfaces 2 and 3, and Δ_b is set to the mean intensity between surfaces 1 and the detected surface 2.

The most important advantage of the multi-resolution approach over the simultaneous segmentation of the the inner surfaces in the full resolution is the considerable reduction in run-time. In the lower resolution, the graph is considerably smaller, thus

Table 4.1: Weights of on-surface cost terms.

Surface	Weight
2	54
3	32.3
4	74
5	99

providing considerably savings in terms of memory requirements as well as processing time. Similarly, in the full resolution, the use of multiple smaller graphs reduces the size of the overall graph.

4.2.3 Determining the Cost Function Parameters

The overall cost function is expressed as a weighted sum of the regional and the on-surface cost terms, but it must be noted that the on-surface costs for each surface can also be weighted in a similar manner to achieve optimal results.

The value of α is determined by testing the segmentation results against the manually edited bronze standard (described in Section 5.2.1), as described by [21], and is it set to the value that minimizes the error. A similar procedure was used to determine the weights of the on-surface cost terms that produced the smallest error on the bronze standard. The values used are tabulated in Table 4.1.

Note, that the value of alpha and the surface cost term weights were determined using a preliminary version of the segmentation method that did not use the multi-resolution approach.

CHAPTER 5
EXPERIMENTAL METHODS

In this chapter we describe the various experimental methods used to validate the methods described in the previous chapter.

5.1 Reference Standard for the Flattening Technique

Tang *et. al* [22] describes a method for the reconstruction of the shape of the optic nerve head from stereo fundus photographs. The 3-D shape estimate is obtained by finding corresponding pixels from two stereo images of the optic nerve head, 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 proportionaly to the depth associated with the 3-D pixel, a disparity map can be created using pixel correspondences. The disparity maps show the shape of the retina at the ONH region, and since the disparity maps are created from fundus photographs, they are free from the curvature artifacts associated with OCT scans. Thus, the disparity maps can be used to evaluate the effectiveness of the flattening procedure.

First, the disparity maps are registered to the OCT images to bring them into the same reference plane; then, the images are smoothed. As the disparity maps also show considerable amounts of noise, in order to retain as much information about the shape of the optic nerve head region, a 2-D thin-plate is fit to the disparity map. Fig. 5.1 shows the resulting smoothed image. The normalized disparity map can now be compared to the flattened ILM to guage the effectiveness of the flattening procedure.

5.2 Reference Standards Used for the Validation of the Segmentation Method

[21] describes the use of two reference standards for the training and testing of the segmentation method. The "bronze" standard was compiled from the 7 training

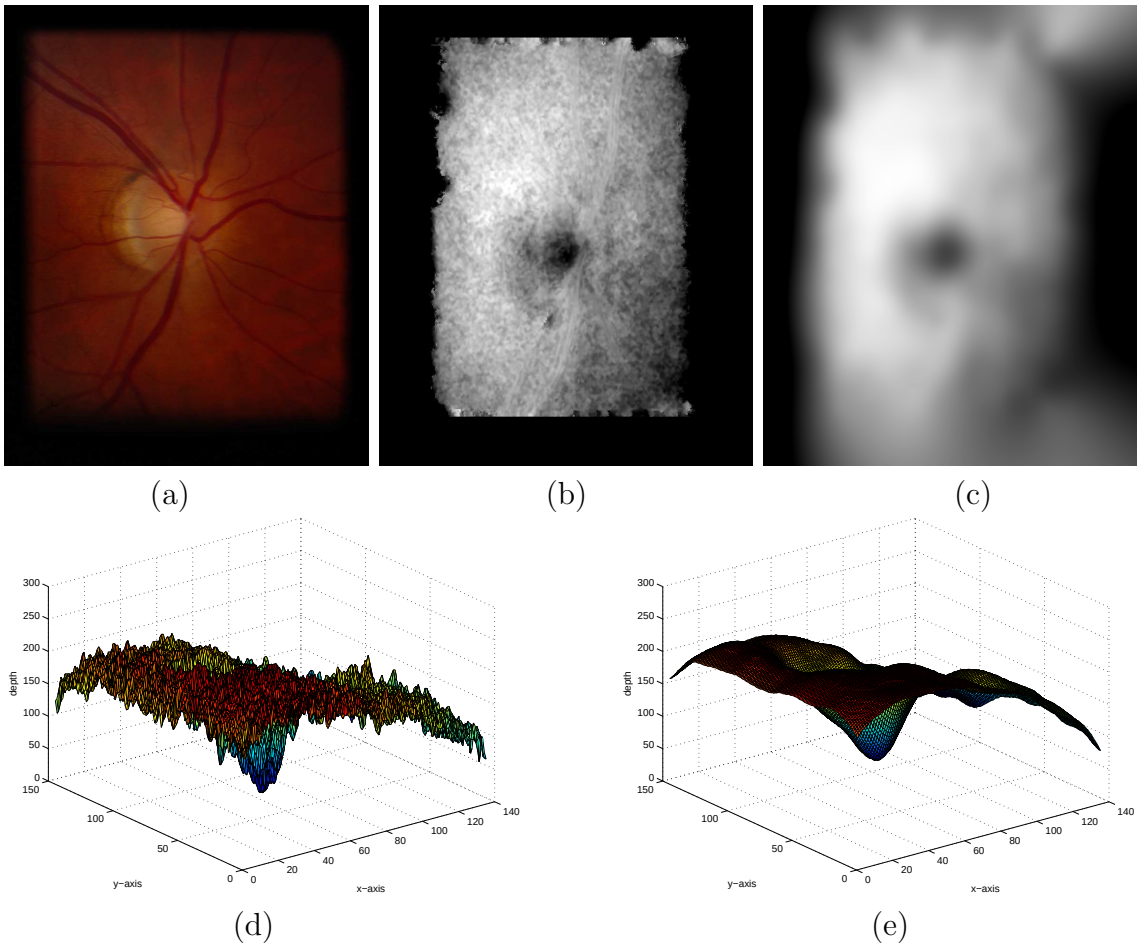

Figure 5.1: Fundus photograph and its corresponding disparity maps. (a) The fundus photograph of a glaucomatous eye. (b) The disparity map constructed from the fundus photographs. (c) The smoothed disparity map obtained after the spline-fit. (d) The disparity map in 3-D. (e) The smoothed disparity map in 3-D showing the overall shape of the optic nerve head.

datasets, and was created from edited manual tracings (from a single observer) of the entire volume. This was done using a special tracing/editing tool and was used to determine the varying feasibility constraints. The "gold" standard consists of manual tracings from 2 independent observers of 10 slices for each testing dataset. The reference standard was then created from the average of the two observer's tracings. Further details regarding each of the reference standard is discussed in the sections below.

Another important aspect that must be considered while validating the segmentation method is the neural canal. Within the neural canal, the surfaces become indistinct and while a human expert may be able to gauge the expected location (but not definitively trace) a surface, it is hard for an automated technique to achieve the same. Thus, this region must be avoided altogether while tracing as well as segmenting the surfaces. Hu *et. al* [23] describes a graph-based method for the automated segmentation of the NCO from OCT datasets. The method begins with the creation of a projection image (mean intensity along the z-axis) from a small number of slices near surface 6. The method also incorporates the location of blood vessels as they are quite large at the ONH. The blood vessels are segmented using the approach developed by Niemeijer *et. al* [24] from the projection image created from the entire flattened dataset. The approach uses a kNN classifier (k=31) and Gaussian filter derivatives up to and including order 2 (i.e. $L, L_x, L_y, L_{xx}, L_{xy}, L_{yy}$) at scales equal to 1, 2, 4, 8, and 16 pixels are used as features. Each pixel is then assigned a soft label to form a posterior map, which can then be thresholded to obtain a map of the large dominant blood vessels (Fig. 5.2(a)). The location of the blood vessels is used to modify the cost function after which, the inner and outer boundaries of the NCO are segmented simultaneously using a graph-search. Fig. 5.2 shows the projection image overlaid with the detected NCO mask.

Once the NCO boundary has been detected, it can be used to avoid trying to validate the performance of the segmentation method within the neural canal. The NCO mask is dilated before being used in the validation process to ensure that no region in the optic canal is included.

5.2.1 Bronze Standard

The bronze standard was created for learning the expected layer thickness and surface smoothness parameters. The slices were traced using a combination of automated segmentation and manual editing. First, the volume was divided into 10 slabs

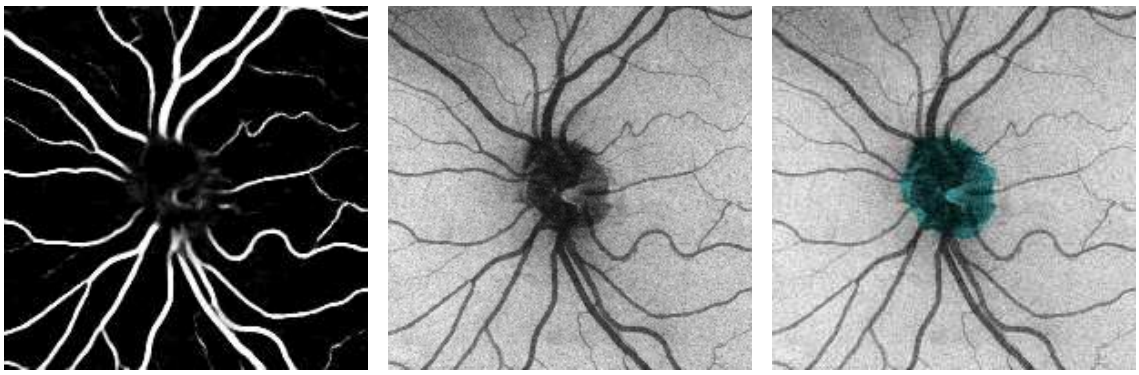

(a) (b) (c)

Figure 5.2: NCO boundary detection. (a) The soft-label vessel map created using the pixel classification method [24]. (b) The projection image created from a small number of slices near Bruch's membrane. (c) The projection image overlaid with the mask of the detected NCO boundary.

and a random slice from each slab was traced manually using the tracing tool. An early version of the segmentation method was then used to segment the surfaces on the entire volume using the traced slices as "anchor" points. The anchoring was achieved by adjusting the cost functions to favor the traced "known" surfaces. The early version of the segmentation method did not use regional terms or varying constraints as the constraint parameters had not been learned yet. Instead, the constraints were set to large enough values to enable the segmentation method to find the surfaces with as much accuracy as possible.

The segmented surfaces were then loaded back into the tracing tool and edited (by a single observer) to form the bronze standard.

5.2.2 Gold Standard

The gold standard, as mentioned earlier, was created for the purposes for validating the segmentation method. Each volume in the dataset was divided into 10 equal sized section, and a slice was chosen at random from each section. These 10 slices were then manually traced by two independent observers to form the gold standard.

The reference standard is formed from the average of the two independent observers' tracings.

5.3 Data Used to Validate the Methods

The images used for the validation were obtained from a spectral-domain OCT scanner (Cirrus, Carl Zeiss Meditec Inc., Dublin, CA) and had dimensions of 6mm× 6mm × 2mm and 200×200×1024 voxels. The test set for the flattening method consisted of 30 OCT scans from 15 patients diagnosed with glaucoma, while the test set for the segmentation method consisted of the gold standard created on 5 normal scans from 5 normal subjects, as well as 10 glaucomatous scans obtained from 10 glaucomatous patients.

CHAPTER 6
VALIDATION AND RESULTS

6.1 Flattening of the Datasets

The flattening method was validated using the disparity maps created as described in Section 5.1. Before the unsigned differences are computed between the disparity maps and the ILM surface, they are normalized to fractional values between 0 and 1. The detected ILM surface - unflattened, flattened using a single spline-fit and flattened using the two-spline approach - are normalized by dividing all the depths by 800 and then scaling the values to within 0 and 1. This is done as the maximum depth of the surface in the unflattened case can exceed the depth of the neural canal, which is the maximum depth seen in the flattened cases.

As the disparity 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 was also avoided. Table 6.1 shows the results obtained on a subset of the test set (only the right eyes), and Fig. 6.1 show some of the flattening results, including the best and worst cases.

Table 6.1 shows the mean unsigned difference between the disparity maps and the ILM in the original unprocessed dataset, the partially flattened dataset (after the 1st stage of flattening) and the final flattened dataset. The mean unsigned difference computed over all 30 datasets datasets flattened using the single-spline and two-spline approach was found to be 0.215 ± 0.056 and 0.129 ± 0.030, ($p < 0.001$) respectively.

6.2 Intraretinal Layer Segmentation

The testing set consisted of 5 datasets from 5 normal subjects and 10 datasets from 10 patients diagnosed with glaucoma. The border positioning errors are expressed as the mean $\pm$ standard deviation expressed over all of the datasets. The results consist of comparisons between the algorithm and each of the observers as well as the

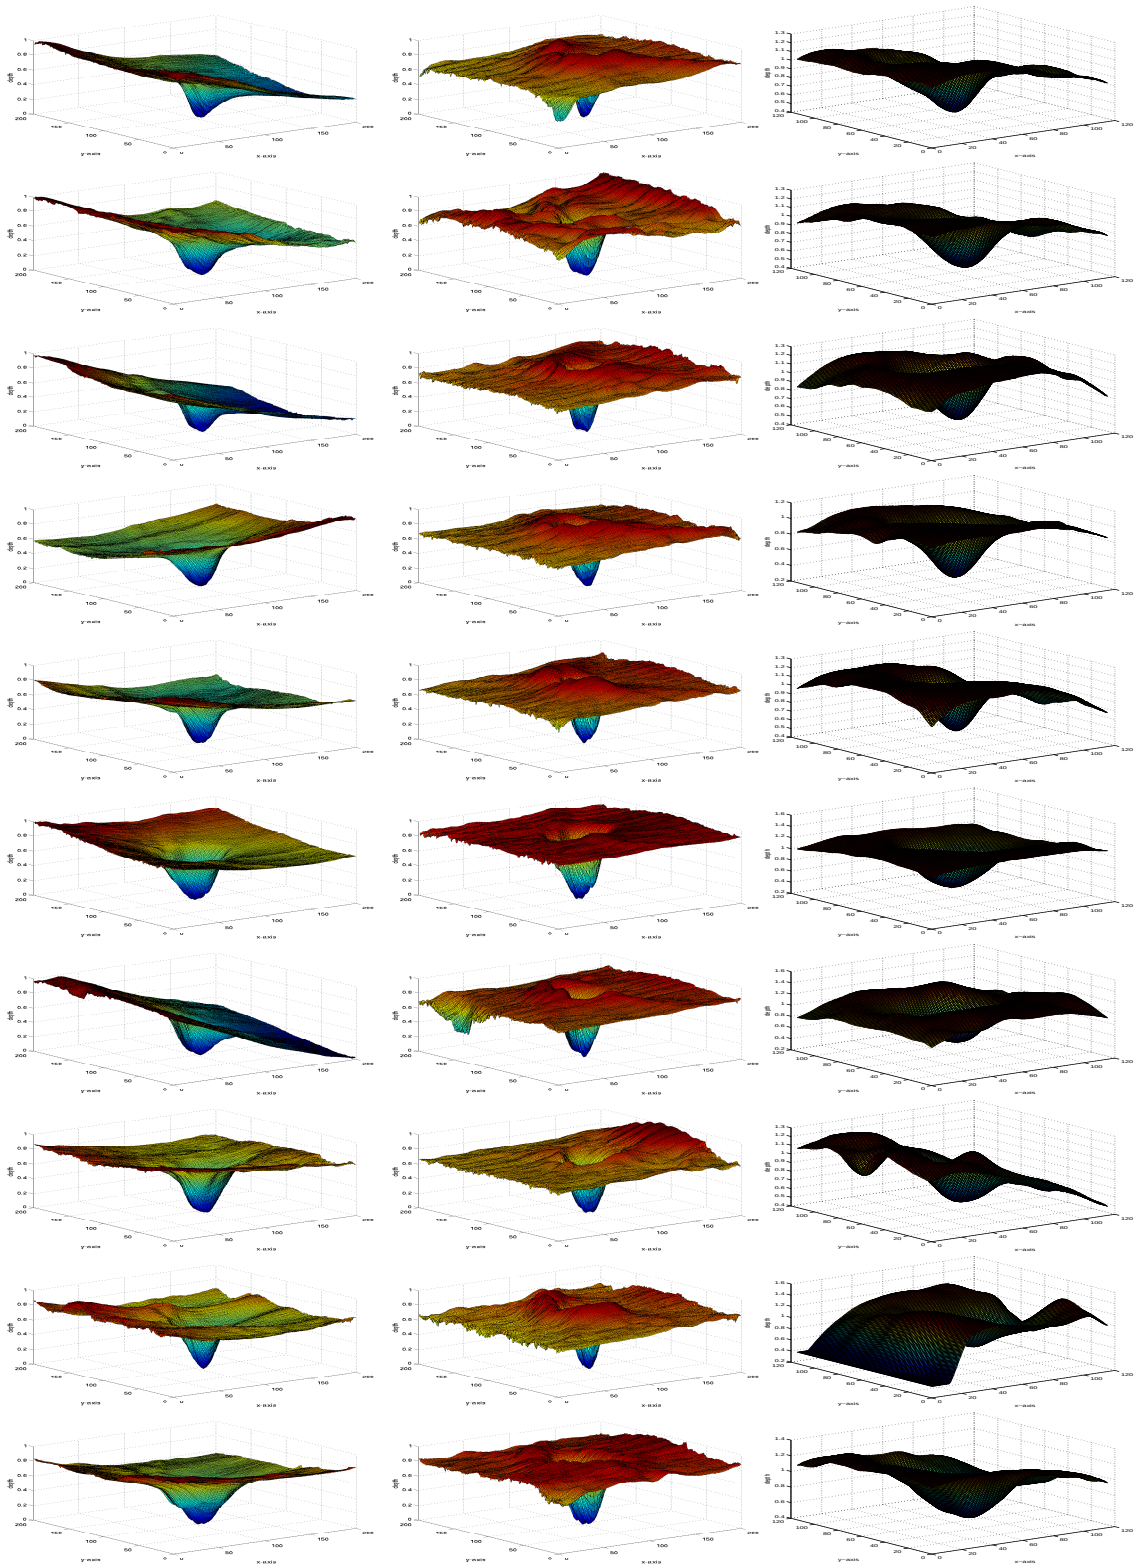

Figure 6.1: Comparison of flattening method. Each row of images show the original segmented ILM, the ILM after the first spline-fit, the ILM after the second spline-fit and the corresponding disparity map created from fundus photographs.

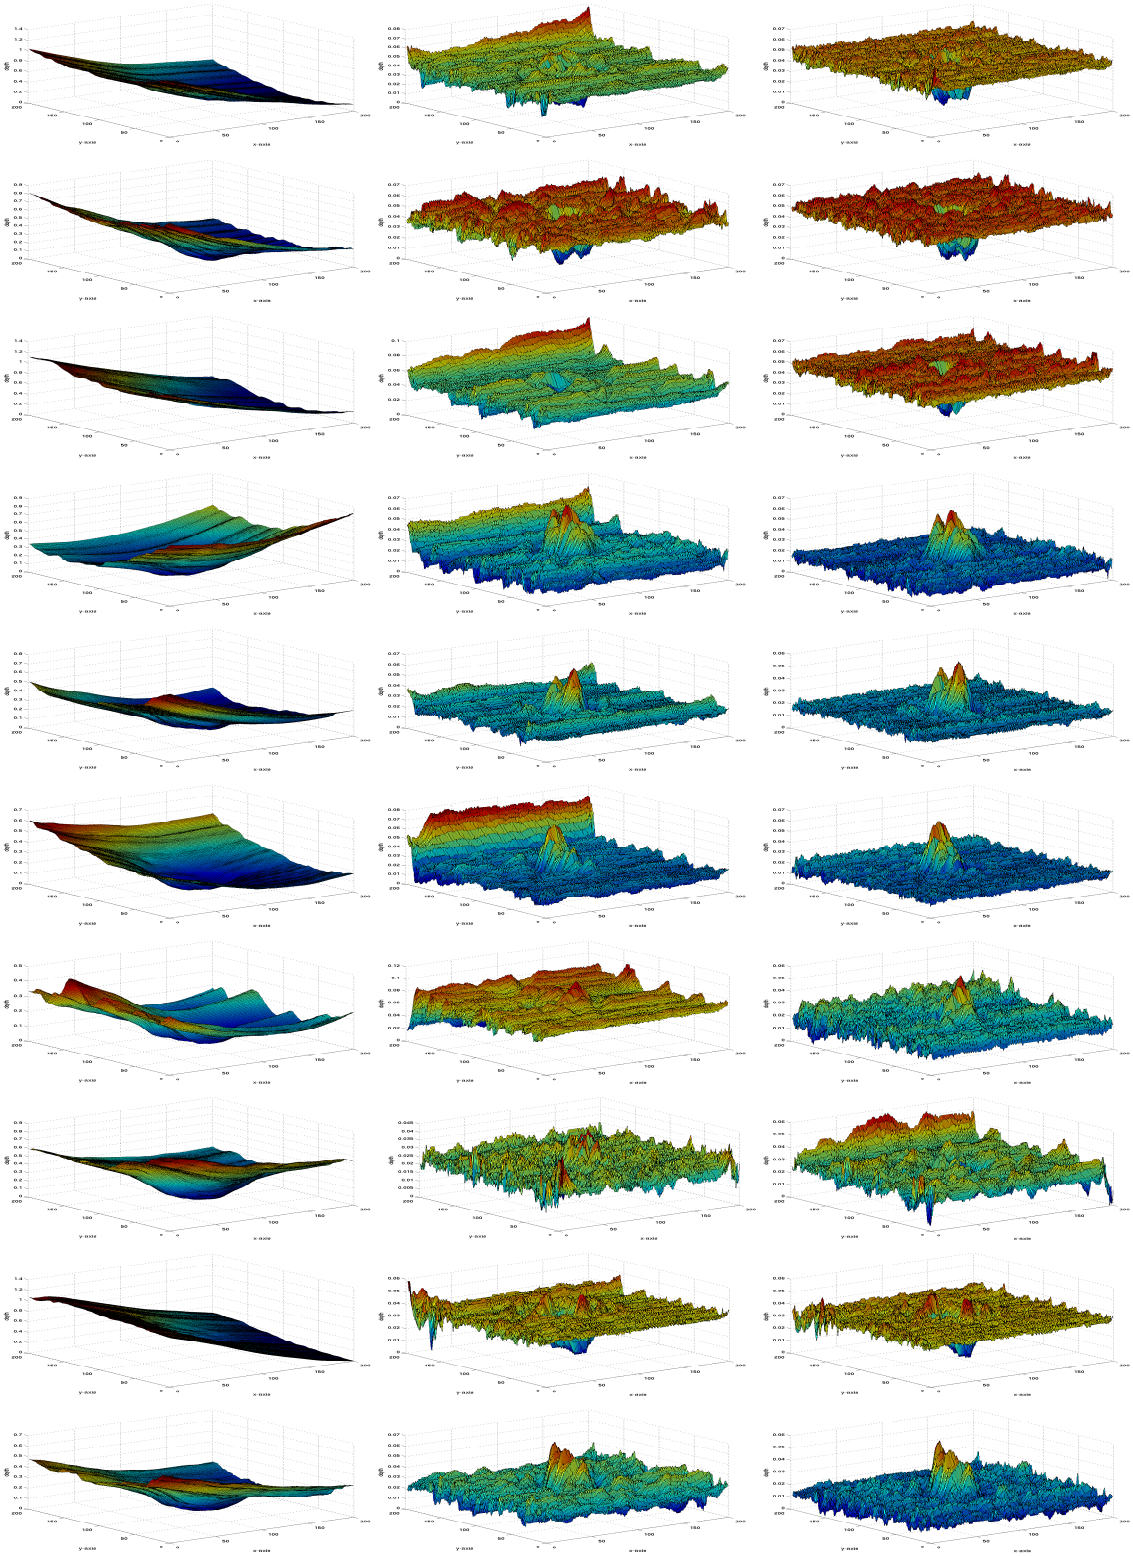

Figure 6.2: Behavior of surface 6 through the flattening process. Each row shows 3-D images of surface 6 before being flattened, flattened after the first spline-fit and flattened after the second spline-fit.

Table 6.1: Mean unsigned difference between the disparity maps and the ILM before and after flattening.*

Dataset	Disparity Map vs. ILM in Original Volume	Disparity Map vs. ILM in Partially Flattened Volume	Disparity Map vs. ILM in Flattened Volume
1	0.196	0.176	0.129
2	0.168	0.156	0.110
3	0.274	0.265	0.133
4	0.296	0.303	0.129
5	0.199	0.202	0.129
6	0.193	0.201	0.099
7	0.059	0.069	0.050
8	0.247	0.259	0.137
9	0.192	0.197	0.119
10	0.231	0.234	0.114
11	0.267	0.264	0.139
12	0.184	0.189	0.128
13	0.156	0.159	0.111
14	0.202	0.197	0.156
15	0.189	0.197	0.149
Average§	0.212 ± 0.059	0.215 ± 0.056	0.129 ± 0.030

* Differences were computed in regions were the disparity maps where well defined. The optic disc was also avoided.

§ Mean ± standard deviation.

the average of the two manual tracings, and the inter-oberser variability. The mean unsigned error of the algorithm when compared against the reference standard was found to be comparable to the inter-oberserver variability, with values of 7.25±1.08 μm and 6.32±1.33 μm, respectively. Table 6.2 shows the mean unsigned errors and Table 6.3 shows the signed errors on the normal test set. The datasets that show the smallest and largest errors are also depicted in Fig. 6.3 and Fig. 6.4, respectively.

The results obtained on the glaucomatous test set are tabulated in Table 6.4 and Table 6.5, and the segmented surfaces on 5 traced slices from the dataset showing the smallest and largest overall errors are shown in Fig. 6.5 and Fig. 6.6 respectively. The performance of the algorithm against the reference standard were again found

to be comparable to the inter-observer variability, and show mean unsigned errors of 8.94 ± 3.76 μm and 8.52 ± 3.61 μm, respectively.

Table 6.2: Mean unsigned border positioning errors† computed on the normal test set.⋆

Surface	Algo. vs. Obs. 1	Algo. vs. Obs. 2	Algo. vs. Avg. Obs.	Obs. 1 vs. Obs. 2
1	4.70±1.59	3.74±0.23	4.09±1.08	4.65± 1.61
2	9.86±1.97	9.38±1.78	10.07±2.16	9.92±2.65
3	9.36 ±1.22	8.21±0.68	9.35 ±1.65	7.48± 1.89
4	8.90 ±1.34	7.85 ±1.03	9.13 ±2.14	7.07 ±1.40
5	10.21±1.52	9.50±1.01	10.04±1.55	7.04 ±1.31
6	2.95±1.58	2.41±0.48	2.89 ±1.05	3.37 ±1.20
7	4.99 ±2.99	5.72±1.73	5.17±1.41	4.74± 2.05
Average	7.28±1.21	6.69 ±0.53	7.25±1.08	6.32±1.33

† Mean ± SD in μm.

⋆ For each boundary, differences were not computed in the neural canal.

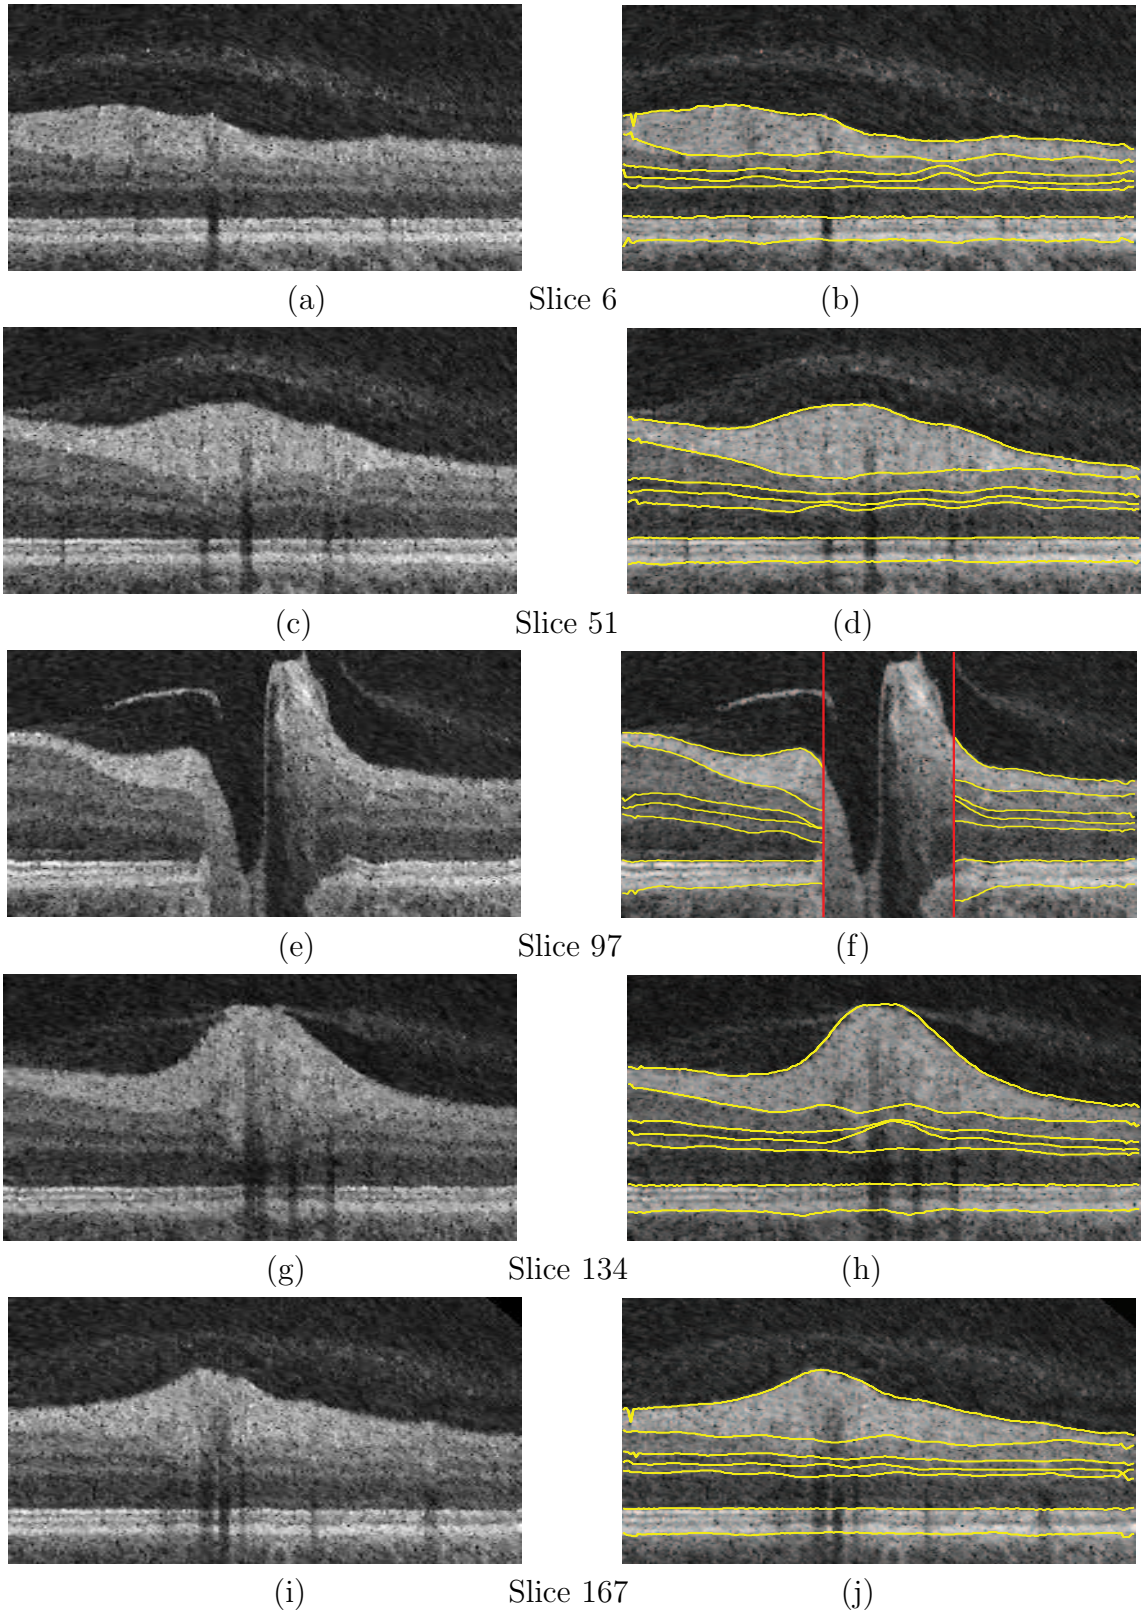

Figure 6.3: Segmented surfaces from a normal test set image that showed the smallest error with respect to the manual tracings of Observer 1. Every other traced slice has been displayed, with the dilated NCO marked in red as applicable. The dataset showed an unsigned error of $5.91 \pm 2.61 \mu$m.

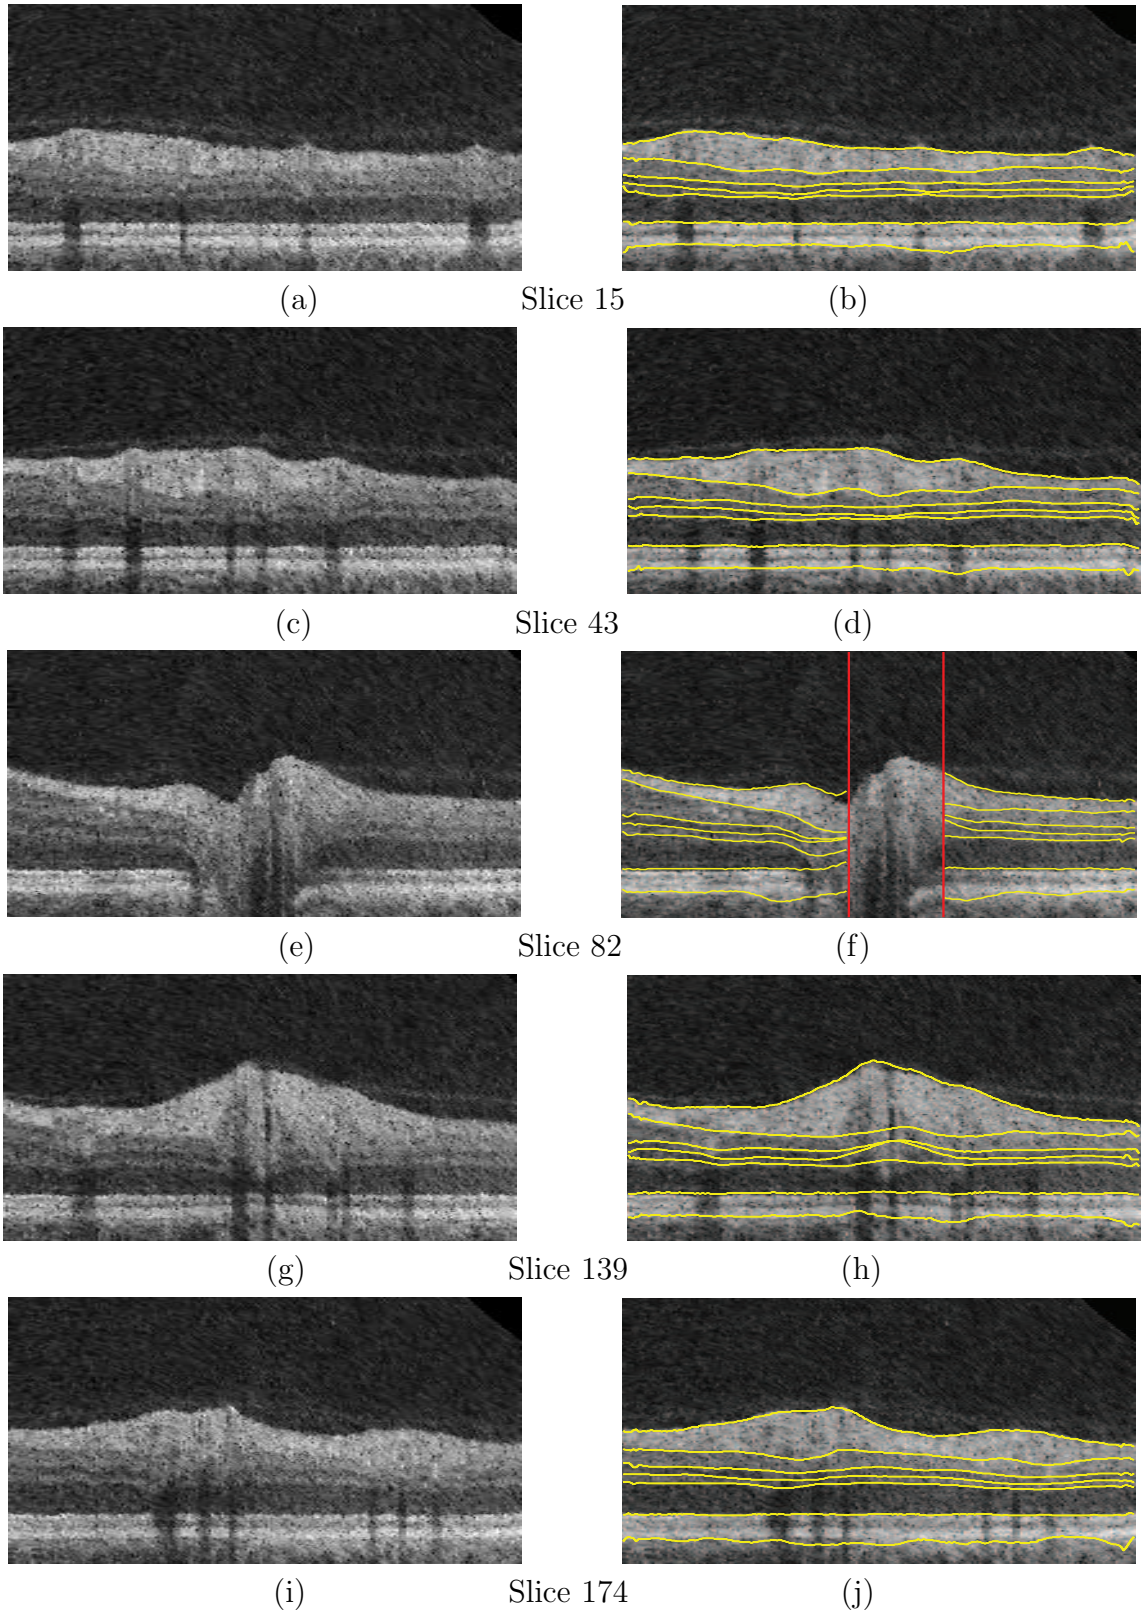

Figure 6.4: Segmented surfaces from a normal test set image that showed the largest error with respect to the manual tracings of Observer 1. Every other traced slice has been displayed, with the dilated NCO marked in red as applicable. The dataset showed an unsigned error of $9.21 \pm 3.16 \mu$m.

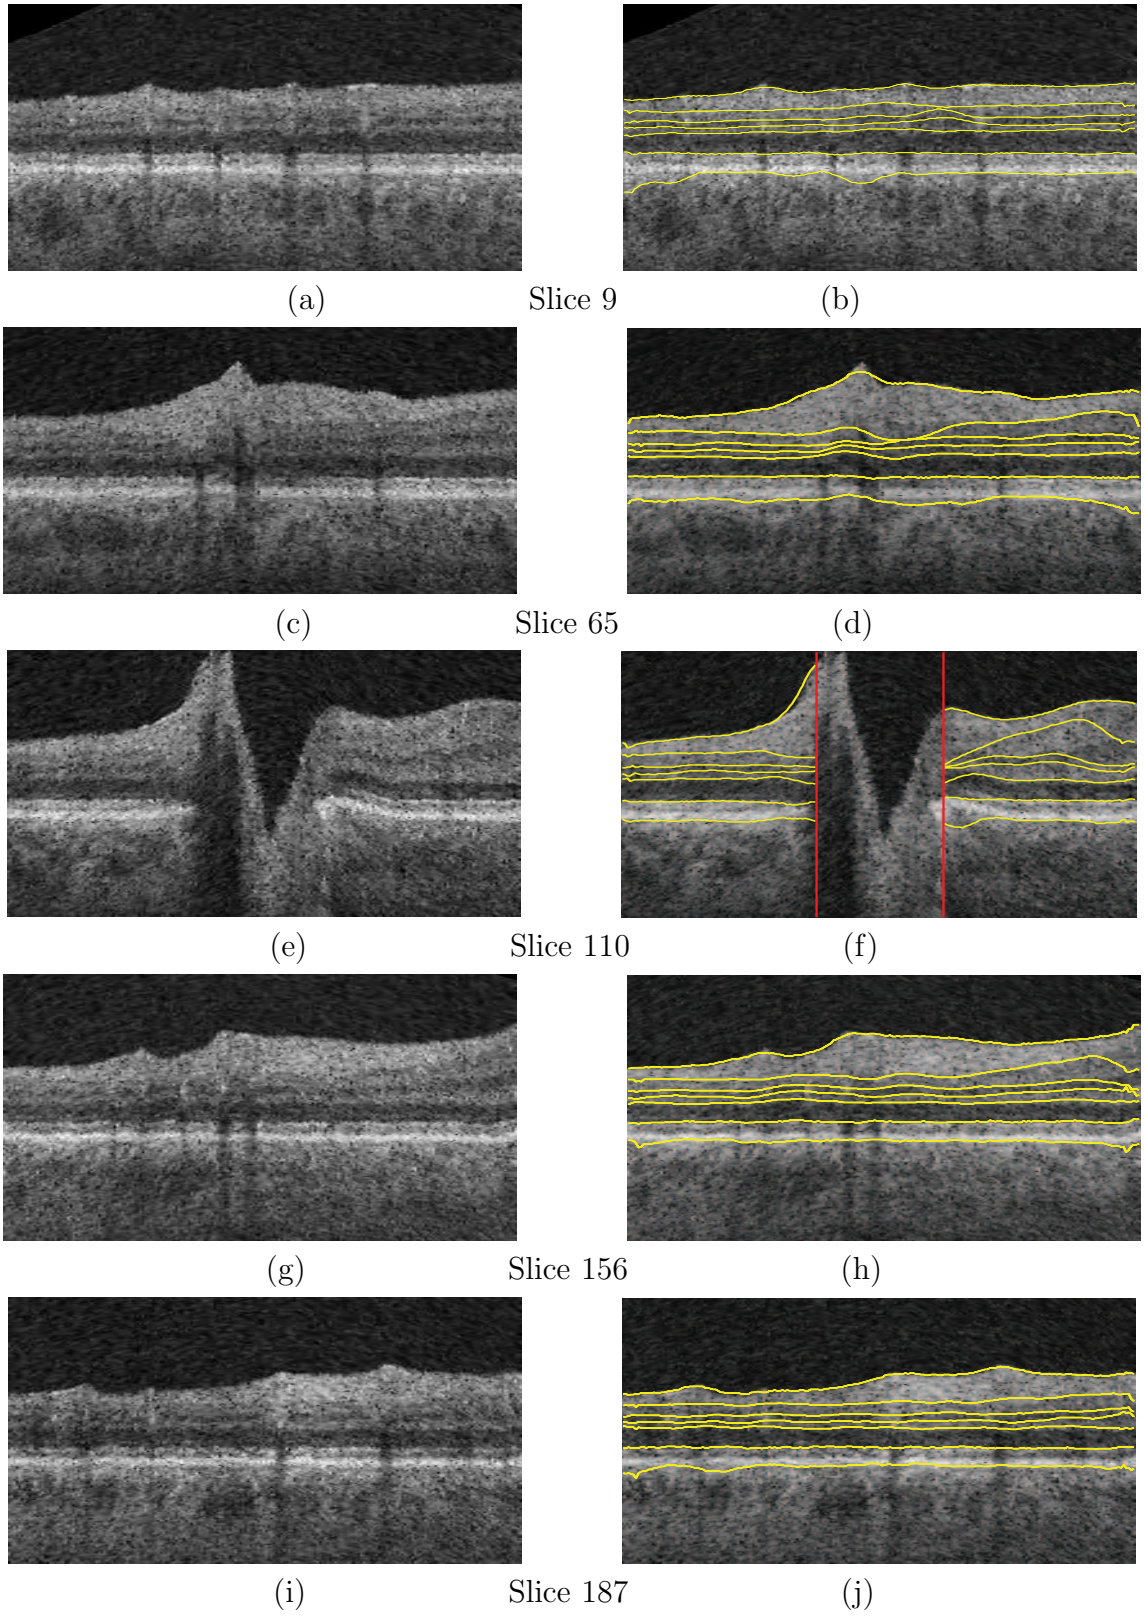

Figure 6.5: Segmented surfaces from a glaucomatous test set image that showed the smallest error with respect to the average manual tracings of both observers. Every other traced slice has been displayed, with the dilated NCO marked in red as applicable. The dataset showed an unsigned error of $6.04 \pm 2.49\mu$m.

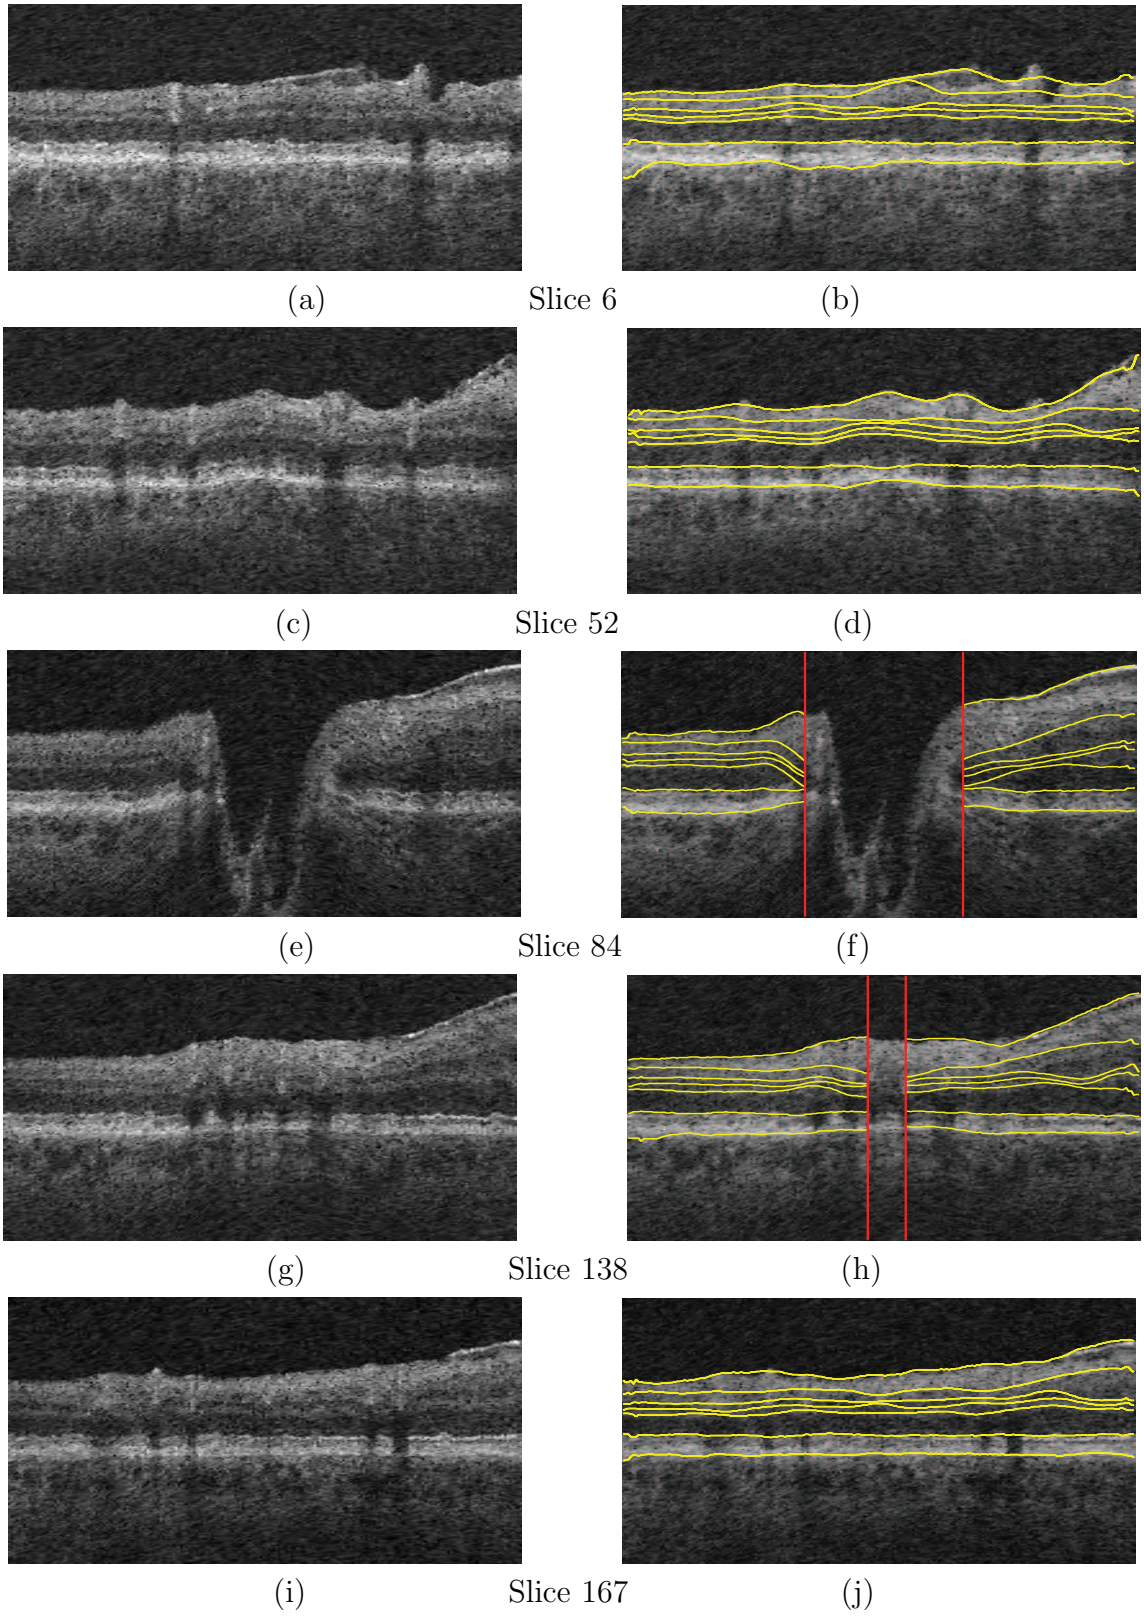

Figure 6.6: Segmented surfaces from a glaucomatous test set image that showed the largest error with respect to the average manual tracings of both observers. Every other traced slice has been displayed, with the dilated NCO marked in red as applicable. The dataset showed an unsigned error of $12.16 \pm 6.52\mu$m.

Table 6.3: Mean signed border positioning errors† computed on the normal test set.⋆

Surface	Algo. vs. Obs. 1	Algo. vs. Obs. 2	Algo. vs. Avg. Obs.	Obs. 1 vs. Obs. 2
1	2.63±0.78	0.57±1.81	1.35±1.92	-1.68±1.79
2	-3.48±1.75	-6.26±1.93	-4.85±3.85	-1.89±2.36
3	-5.75±1.19	-3.99±0.81	-5.63±0.78	2.23 ±1.68
4	-5.39±1.71	-4.40±2.12	-6.20±1.76	0.96 ±2.10
5	-8.36±0.98	-8.80±1.49	-8.91±1.69	0.26 ±2.45
6	0.48 ±0.52	-0.34±1.09	-0.47±1.09	-1.03±1.42
7	-0.69±2.45	-2.34±3.56	-1.39±2.62	-1.25±3.62

† Mean ± SD in μm.

⋆ For each boundary, differences were not computed in the neural canal.

Table 6.4: Mean unsigned border positioning errors† computed on the glaucomatous test set.⋆

Surface	Algo. vs. Obs. 1	Algo. vs. Obs. 2	Algo. vs. Avg. Obs.	Obs. 1 vs. Obs. 2
1	5.26 ±1.61	4.66 ± 1.58	4.90 ± 1.54	4.90 ±1.37
2	14.29 ± 6.28	15.12 ± 4.44	14.43 ± 5.63	12.79 ± 3.36
3	10.32 ± 2.75	14.17 ± 4.98	10.96 ± 4.06	13.74 ± 2.04
4	10.77 ± 3.37	10.44 ± 3.06	10.46 ± 2.79	9.28 ± 3.00
5	10.27 ± 2.82	11.04 ± 3.05	10.73 ± 2.78	7.67 ± 1.69
6	3.73 ± 1.26	4.07 ± 1.30	3.87 ± 1.32	4.69 ± 1.26
7	7.73 ± 1.59	6.98 ± 1.75	7.24 ± 1.74	6.58 ± 1.53
Average	8.94 ± 1.97	9.50 ± 4.39	8.94 ± 3.76	8.52 ± 3.61

† Mean ± SD in μm.

⋆ For each boundary, differences were not computed in the neural canal.

Table 6.5: Mean signed border positioning errors† computed on the glaucomatous test set.⋆

Surface	Algo. vs. Obs. 1	Algo. vs. Obs. 2	Algo. vs. Avg. Obs.	Obs. 1 vs. Obs. 2
1	2.56 ± 2.61	2.29 ± 1.69	2.80 ± 1.73	0.04 ± 2.19
2	-3.73 ± 1.89	6.28 ± 6.02	7.26 ± 5.78	0.12 ± 4.47
3	-5.94 ± 1.10	10.18 ± 6.02	4.90 ± 5.61	-11.33 ± 2.12
4	-5.221 ± 1.51	-2.01 ± 4.03	-2.37 ± 4.14	-0.80 ± 2.80
5	-8.11 ± 0.99	-5.50 ± 2.93	-5.46 ± 3.21	-0.18 ± 1.79
6	0.01 ± 1.16	1.29 ± 1.60	1.19 ± 1.47	-0.18 ± 1.79
7	-0.96 ± 2.50	1.58 ± 1.84	1.16 ± 2.90	-0.41 ± 2.80

† Mean ± SD in μm.

⋆ For each boundary, differences were not computed in the neural canal.

CHAPTER 7
POSSIBLE CLINICAL APPLICATIONS

Glaucoma is known to thin the RNFL as it progresses, and the degree of thinning (currently available on a 2-D circular scan) is often used for the diagnosis of the disease. To demonstrate a possible application of this segmentation method, we have tabulated the layer thickness values (of the RNFL and other layers) on a set of 26 normals and 70 glaucomatous datasets. Table 7.1 and Fig. 7.1 show the mean layer thickness values computed on all 6 layers that were segmented. The RNFL, as expected, shows a marked reduction in the thickness with average values of 73.72 $\pm$ 32.72 μm in normals and 60.38 $\pm$ 25.22 μm in glaucomatous eyes. The optic disc region has not been included in the evaluation of the thickness maps as the surfaces become indistinct and the layers are hard to define in this region. The area has been approximated using a circular mask.

Table 7.1: Average Layer thicknesses[†] in normal and glaucomatous eyes.

Layer	Normals	Glaucomatous
1	73.72 $\pm$ 32.72	60.38$\pm$ 25.22
2	44.61$\pm$ 14.79	42.37$\pm$ 11.21
3	19.40$\pm$ 2.04	17.74 $\pm$2.24
4	20.32$\pm$ 4.56	20.10 $\pm$4.79
5	74.10$\pm$ 6.03	69.48$\pm$ 7.80
6	59.15$\pm$ 2.50	58.25$\pm$ 3.53

[†] Thickness in mean $\pm$ standard deviation μm.

[⋆] Average was computed from 26 normal and 70 glaucomatous datasets.

The thinning of the RNFL is also known to be localized and thus, there is known to be some advantages to calculating the mean thicknesses in sectors instead of the entire

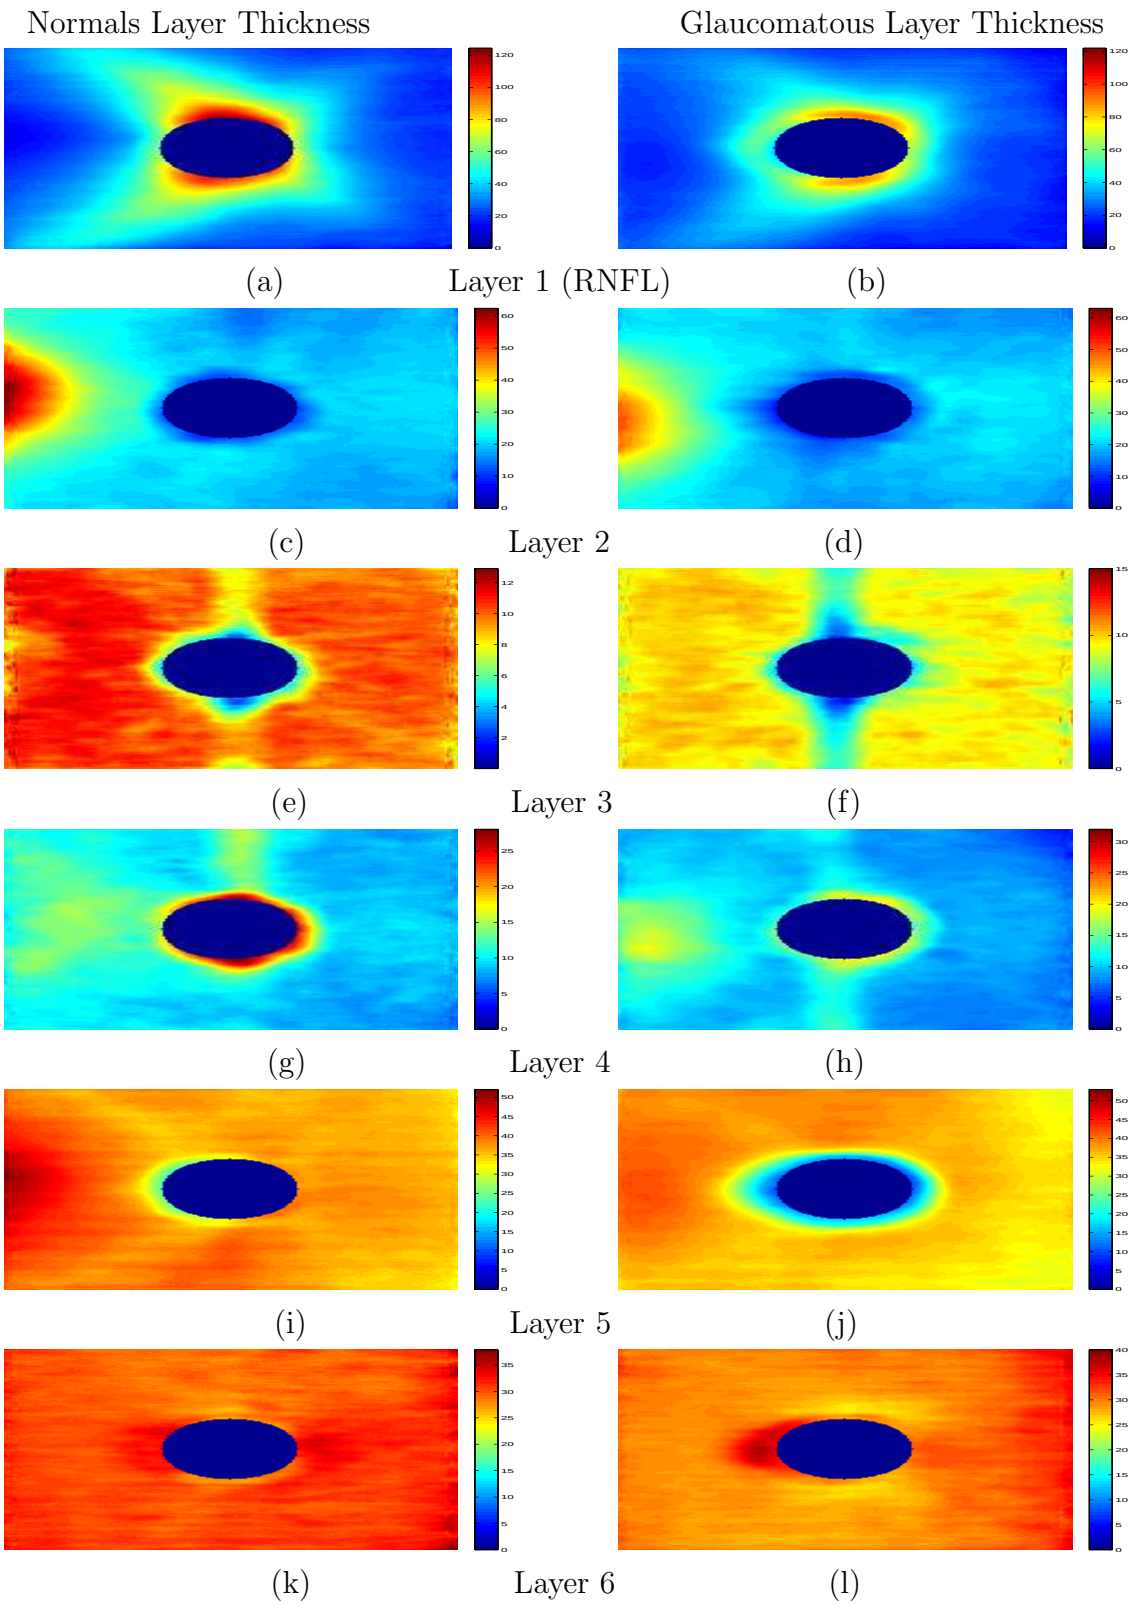

Figure 7.1: Mean thickness maps of normal and glaucomatous eyes.

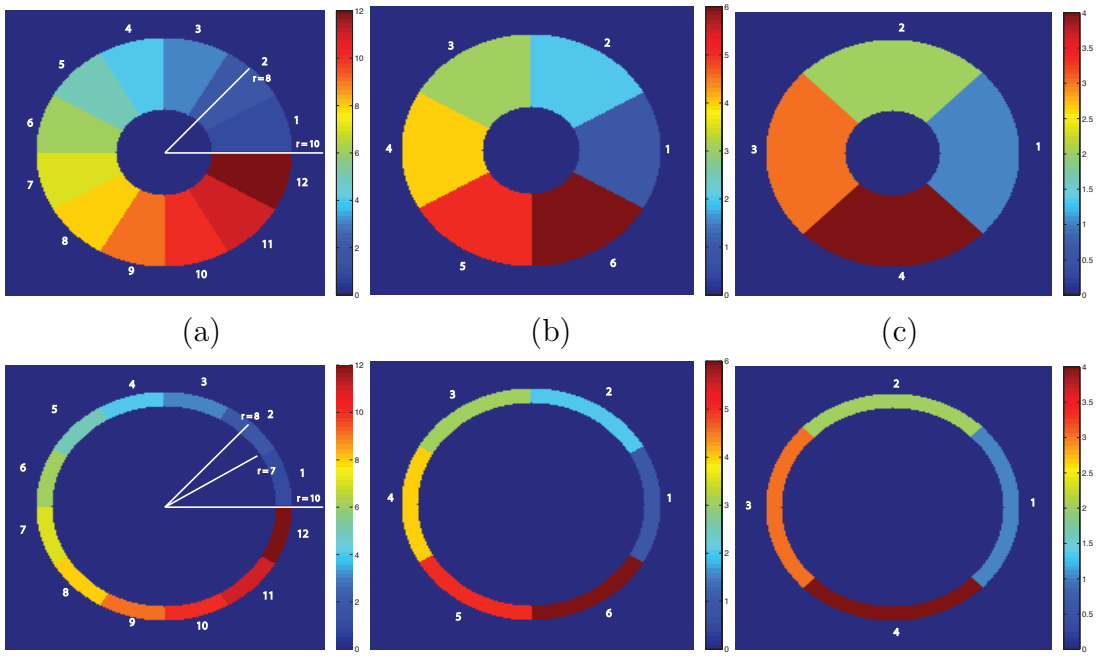

Figure 7.2: Protocols used in the computation of mean RNFL thickness values.

volume. Fig. 7.2 shows the different methods that were used to divide the image into sectors. In addition to the neural canal, the peripheral regions were also excluded as it is the nerve fibers closer to the optic disc that show thinning in glaucomatous eyes. Apart from this, we also evaluated mean thickness values in a narrow "doughnut" shaped region, within two radial distances. If the largest circle that can be drawn in the 200×200 pixel image is assumed to have a radius of 10, then the doughnut-shaped region is defined between radial distances of 7 and 8.

The graphs shown in Fig. 7.3 show the average RNFL thicknesses computed within the narrow doughnut-shaped region, while dividing the image into 4, 6 and 12 sectors. The plot shows us the sectors that show the largest difference in layer thicknesses (thinning) between normal and glaucomatous scans.

Table 7.2: Average RNFL thicknesses † in normal and glaucomatous eyes in 12 equal sectors. ⋆

	Sector within r=8		Sector within r1=7, r2=8	
Sector	Normals	Glaucomatous	Normals	Glaucomatous
1	69.38± 25.03	69.79± 26.63	47.93± 2.54	49.62± 4.12
2	95.45± 34.11	88.17± 32.31	65.98± 7.56	59.94± 4.51
3	110.34± 43.27	89.72± 34.34	70.8± 6.61	60.08± 2.74
4	138.76± 33.89	96.76± 27.36	107.31± 14.06	72.29± 6.84
5	104.26 ±29.34	82.89± 22.3	83.77± 17.76	63.24± 7.12
6	60.11± 20.81	68.67± 22.61	42.22± 6.14	46.3± 3.89
7	71.75± 22.39	62.84± 21.63	53.18± 8.48	41.01± 2.89
8	113.28± 27.26	81.59± 19.71	93.38± 15.44	64.87± 10.83
9	124.77± 34.12	99.59± 26.95	93.13± 15.92	72.67± 6.76
10	98.49± 36.54	85.13± 36.68	64.51± 4.34	51.83± 3.83
11	102.69± 30	75.07± 29.75	74.17± 6.09	49.54± 3.03
12	76.53 ±27.2	62.7± 24.68	53.91± 5.94	43.23± 1.09

† Thickness in mean ± standard deviation μm.

⋆ Average was computed from 26 normal and 70 glaucomatous datasets.

Table 7.3: Average RNFL thicknesses † in normal and glaucomatous eyes in 6 equal sectors. ⋆

	Sector within r=8		Sector within r1=7, r2=8	
Sector	Normals	Glaucomatous	Normals	Glaucomatous
1	73.02± 26.39	66.18± 25.89	50.95± 5.47	46.39± 4.38
2	102.95± 39.70	88.95± 33.35	68.39± 7.50	60.01± 3.73
3	121.35± 36.06	89.76± 25.88	95.39± 19.89	67.71± 8.32
4	65.83± 22.37	65.81± 22.32	47.63± 9.19	43.69± 4.33
5	119.07± 31.43	90.66± 25.29	93.25± 15.67	68.77± 9.84
6	100.61± 33.46	80.05± 33.73	69.40± 7.17	50.67± 3.63

† Thickness in mean ± standard deviation μm.

⋆ Average was computed from 26 normal and 70 glaucomatous datasets.

Table 7.4: Average RNFL thicknesses † in normal and glaucomatous eyes in 4 equal quadrants. ⋆

	Sector within r=8		Sector within r1=7, r2=8	
Sector	Normals	Glaucomatous	Normals	Glaucomatous
1	79.84± 29.67	70.26± 28.11	55.90± 9.47	48.55± 5.69
2	120.03± 38.90	92.15± 30.11	87.72± 19.64	66.05± 7.10
3	75.33± 26.69	68.88± 22.02	56.66± 15.70	47.98± 7.56
4	113.38± 35.13	89.71± 31.05	83.25± 18.82	62.39± 12.04

† Thickness in mean ± standard deviation μm.

⋆ Average was computed from 26 normal and 70 glaucomatous datasets.

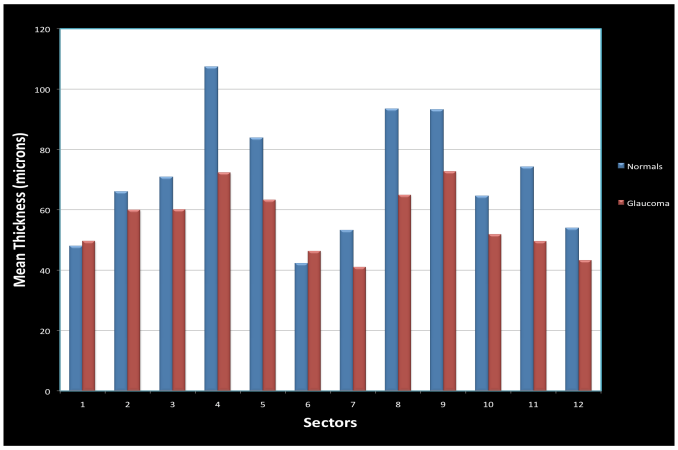

(a)

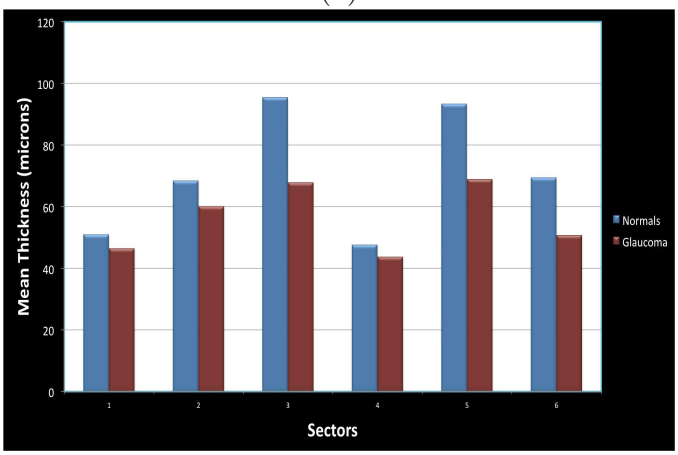

(b)

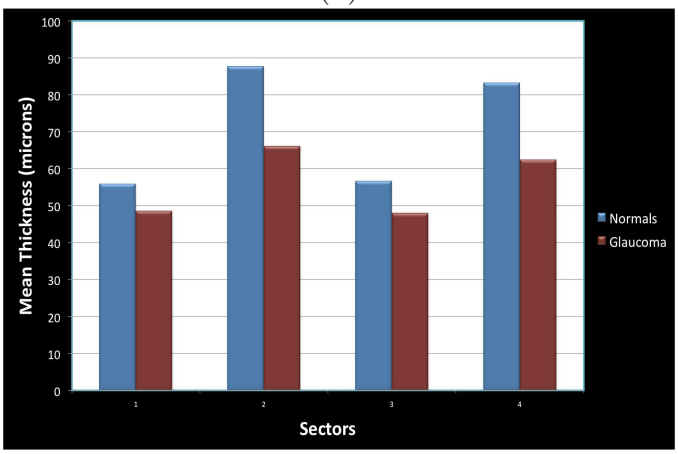

(c)

Figure 7.3: Graph of normal and glaucomatous RNFL thicknesses computed between r1=7 , r2=8. (a) Shows the plot of the mean values computed in the 12 sectors. (b) Shows the plot of the mean values computed in the 6 sectors. (c) Shows the plot of the mean values computed in the 4 sectors.

CHAPTER 8
DISCUSSION

This thesis presents an automated 3-D graph-based method for the segmentation of 7 surfaces (6 layers) from OCT scans centered on the ONH. In addition, we also describe a new method for flattening the OCT datasets and provide some possible clinical applications of the segmentation method.

The flattening method, as described earlier, uses two levels of spline-fits in order to compensate for the two distinct artifacts seen in OCT scans. While it is possible to attain the same results using a large number of control points in a single spline-fit, the memory and time requirements of such an approach makes it less practical. The use of two spline-fits also allows us to use larger values for the spline regularization which makes the method less dependent on the segmentation results. But it is not possible to completely eliminate this dependence and the accuracy of the segmentation still plays a big role in the result. Similarly, the validation is also dependent on the accuracy of the disparity maps, which are prone to noise. Another aspect that affects the validation is the size of the region that the disparity maps reconstructs, which is limited to a small region around the optic nerve head, while the OCT scans image a much larger region. This limits the area that can be used to validate the flattening process, and it is often limited to a narrow band on either side of the optic nerve head.

The segmentation method proposed incorporates two key aspects of two 3-D graph-based approaches that have been proposed for the segmentation of surfaces at the macula. Garvin *et. al* [3] incorporated regional information in the cost function, and used varying feasibility constraints that were derived from a training set. But this method shows large run-times and memory requirements. To tackle these issues, Lee [4] proposed a multi-resolution approach that used several graph-searches to find the surfaces instead of detecting them simultaneously. This method does not in-

corporate regional information or varying feasibility constraints, and boasts run-times of a few minutes. The method proposed in this thesis uses varying constraints and regional information, as well as the multi-resolution approach. The method is robust and is capable of providing comparable results on normal as well as diseased cases. The most common errors seen in the normal test set were on surfaces 3, 4 and 5, and it is likely that the optimization of the cost function parameters will provide a better solution. Also, the errors seen in surface 7 could possibly be fixed by the inclusion of regional information and smaller learned constraints. Another aspect that must be considered is the learned constraints itself. Although the thickness constraints are not affected by flattening, the smoothness constraints (especially in the y-direction) are likely to show very large variations which are no longer seen in the datasets that have been flattened with the new technique. Thus, further editing of the training set to reflect the likely y-axis smoothness might prove useful. The large blood vessels at the optic disc also pose quite a challenge and while the method is capable of dealing with the smaller vessel shadows, it fails within regions where multiple vessel shadows occur close together. The incorporation of vessel information, possibly to adjust the cost function or the feasibility constraints, might also help reduce such errors.

The glaucoma scans show much larger errors, due to a combination of the above mentioned factors as well as low contrast in the images itself. This is reflected in the inter-oberserver variability, which reflects the difficulty faced by the observers while tracing the scans manually. The inclusion of glaucomatous scans in the training set could prove useful, as the training set would now reflect scans with strong edge information as well as scans that typically show poor contrast. The incorporation of texture information in addition to the regional information might also help make the method more robust, as the texture information would not be entirely dependent on intensity values.

CHAPTER 9
CONCLUSION

We have presented an automated 3-D method for the segmentation of intraretinal layers from spectral-domain OCT scans centered on the ONH. We have also described and quantitatively validated an alternative flattening approach that address the artifacts common to OCT datasets. Although various methods for the same have been proposed, thus far, none have validated the results quantitatively.

The flattening technique addresses the two distinct artifacts seen in OCT datasets, and the advantages of this method over other flattening methods has been demonstrated by comparing the results with 3-D reconstructions of the optic nerve head.

The intraretinal surface segmentation method addresses the challenges posed by the presence of the optic disc and the large blood vessels found at this region of the retina. It is sufficiently robust to deal with normal as well as diseased scans (which often show low contrast), and produced border-positioning errors that are comparable with inter-observer variability.

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] K. Li, X. Wu, D. Z. Chen, and M. Sonka, "Optimal surface segmentation in volumetric images – a graph-theoretic approach," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 1, pp. 119–134, 2006.

[3] 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 Transactions on Medical Imaging*, vol. 28, no. 9, pp. 1436–47, 2009.

[4] K. M. Lee, "Segmentations of the intraretinal surfaces, optic disc and retinal blood vessels in 3-D OCT scans," Ph.D. dissertation, The University of Iowa, 2009.

[5] B. Klein, R. Klein, W. Sponsel, T. Franke, L. Cantor, J. Martone, and M. Menage, "Prevalence of glaucoma. the beaver dam eye study." *Ophthalmology*, vol. 99, no. 10, p. 1499, 1992.

[6] D. Friedman, R. Wolfs, B. O'colmain, B. Klein, H. Taylor, S. West, M. Leske, P. Mitchell, N. Congdon, and J. Kempen, "Prevalence of open-angle glaucoma among adults in the united states," *Archives of Ophthalmology*, vol. 122, no. 4, p. 532, 2004.

[7] 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.

[8] 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

[9] 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.

[10] F. Medeiros, L. Zangwill, C. Bowd, R. Vessani, R. S. 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," *American Journal of Ophthalmology*, vol. 139, no. 1, pp. 44–55, 2005.

[11] A. Dichtl, J. Jonas, and G. Naumann, "Retinal nerve fiber layer thickness in human eyes," *Graefe's Archive for Clinical and Experimental Ophthalmology*, vol. 237, no. 6, pp. 474–479, 1999.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] 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*, pp. 679–687, Nov 2008.

[17] X. Wu and D. Z. Chen, "Optimal net surface problems with applications," in *Proceedings of the 29th International Colloquium on Automata, Languages, and Programming (ICALP), LNCS 2380*. Springer-Verlag, 2002, pp. 1029–1042.

[18] 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 *Proceedings of the 10th International 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.

[19] M. Haeker, 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 *Information*

Processing in Medical Imaging (IPMI), ser. Lecture Notes in Computer Science, vol. 4584. Springer, 2007, pp. 607–618.

[20] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion." *IEEE Transactions on Pattern Analysis Machine Intelligence*, vol. 12, no. 7, pp. 629–639, 1990.

[21] M. K. Garvin, "Automated 3-d segmentation and analysis of retinal optical coherence tomography images," Ph.D. dissertation, The University of Iowa, 2008.

[22] L. Tang, Y. H. Kwon, W. L. M. Alward, E. C. Greenlee, K. Lee, M. K. Garvin, and M. D. Abràmoff, "3-D reconstruction of the optic nerve head using stereo fundus images for computer-aided diagnosis of glaucoma," accepted for publication, SPIE 2010.

[23] Z. Hu, M. Niemeijer, K. Lee, M. D. Abràmoff, M. Sonka, and M. K. Garvin, "Automated segmentation of the optic disc margin in 3-d optical coherence tomography images using a graph-theoretic approach," in *Proceedings of SPIE Medical Imaging 2009: Biomedical Applications in Molecular, Structural, and Functional Imaging*, 2009, pp. 7262–29.

[24] M. Niemeijer, M. K. Garvin, B. van Ginneken, M. Sonka, and M. D. Abràmoff, "Vessel segmentation in 3D spectral OCT scans of the retina," in *Proceedings of SPIE Medical Imaging 2008: Image Processing*, vol. 6914. SPIE, 2008.