

Quantifying Atherosclerotic Changes
with Digital Image Processing

Konstantinos Chandrinos

MSc Information Technology: Knowledge Based Systems

Intelligent Robotics Theme

Department of Artificial Intelligence

University of Edinburgh

September 1995

Αφιερώνεται στον πατέρα μου Βασιλη Χαυδρινο

Abstract

The aim of this project was to design and implement a non-invasive, automated method for the examination and follow up of the arteriosclerotic changes due to hypertension, with the help of digital image processing of the retinal vessels. This would help a lot in evaluating the efficacy of various treatments on the regression and reversion of arteriosclerotic lesions. We develop a method for segmenting vessels from fundus images using ridge detection. An algorithm is devised and implemented. Once the vessels are segmented we move on to measurements of vessel width, thickness of the wall, length of the vessels and tortuosity index. All our methods are automatic and do not depend on any image specific thresholds or other parameters.

Acknowledgements

It is a lot of people I would like to thank because without them this project would never have come to an end. I would first of all like to thank my supervisor, Dr Fisher, for his good guidance throughout my research period. George Kyriazopoulos M.D. who actually proposed this project, helped me a lot with the medical side of the problem. Maurizio Pilu (PhD student) was my code angel when I thought everything was going down the drain. I would also like to thank Andrew Fitzgibbon and the rest of the Machine Vision Unit researchers that supported me with code and advice. Douglas Howie from the Departmental lab was kind enough to share with me some of his vast knowledge on film and photographic techniques. Allan James Frame from University of Aberdeen pointed me towards some very interesting papers. And, of course, thank God for Elena: if it hadn't been for her, I wouldn't have been me....

Table of Contents

1. Introduction	1
1.1 Computers in Medicine	2
1.2 Overview of the Project	3
1.2.1 Objectives	4
1.2.2 Achievements	4
1.2.3 Conclusion	5
1.3 A Guide to the Rest of the Thesis	5
2. Background	7
2.1 Medical Background	7
2.1.1 Vessel Calibre	8
2.1.2 Heamodynamics	9
2.2 Image Processing background	10
2.2.1 Fundus Photography	10
2.2.2 Previous Work on Retinal Images	11
2.3 Images Used in this Study	12
2.3.1 Image Acquisition	12
2.3.2 Digitization	13
2.3.3 Image Formats	13
2.3.4 Software Implementation	14
2.3.5 Our Demonstration Image	14
2.4 Discussion	15
3. Calibration	16
3.1 Papilla Fitting	17
3.2 Experimental Results	19
3.3 Discussion	20

4. Extraction of retinal vessels	22
4.1 Initial Approaches	23
4.1.1 Thresholding	23
4.1.2 Edge Detecting operators	23
4.2 Ridge Detection	25
4.2.1 Smoothing	27
4.2.2 Non-maximal suppression	31
4.2.3 Local and Contrast Thresholding	32
4.2.4 Preparation for Tracking	33
4.3 Tracking	34
4.4 Discussion	38
5. Measurements of Retinal Structures	42
5.1 Arteriolar length	42
5.2 Estimation of Vessel Calibre and Wall	43
5.2.1 Sub-pixel Precision	48
5.3 Tortuosity	49
5.4 Discussion	51
6. Conclusions	55
6.1 Summary of goals	55
6.2 Summary of the work completed	56
6.3 Evaluation of the work	56
6.4 Future work	57
Appendices	
Appendices	60
A. Experimental Results	61

List of Figures

1-1	The human eye: A wealth of information	3
2-1	The image used for demonstration of our methods throughout this document is a 400x430 B/W fundus photograph with the papilla centred	14
3-1	Contrast enhancement between two areas that the mask looks at . .	19
3-2	Papilla location with the first method proposed	19
3-3	Papilla location with the second method proposed	20
3-4	Papilla location in a different image with the second method	20
4-1	A histogram of our image	23
4-2	An attempt to threshold	24
4-3	Results of Canny edge detector	25
4-4	A vessel profile before smoothing	26
4-5	Vessel of figure 4-4 after smoothing	26
4-6	The steps of our method for vessel extraction	28
4-7	A typical vessel, forming clearly a valley of intensity values	29
4-8	Vessel from Figure 4-7 smoothed. A ridge has now taken the place of the valley	29
4-9	A horizontal vessel before smoothing	29
4-10	The vessel from Figure 4-9, after smoothing	30

4-11	The Gaussian is not confused at vessel bifurcations	30
4-12	Ridge direction map	32
4-13	Non-maximal Suppression	33
4-14	Our image thinned and ready for tracking	35
4-15	The final tracks are superimposed on the original grey level image. .	36
4-16	A magnified part of three tracked vessels	37
4-17	Tracking results in a different image	37
4-18	Detail of an arteriovenous crossing	38
4-19	A fraction of our test image	39
4-20	Different levels of smoothing on a vessel	40
4-21	Linking by interpolation	41
5-1	Application of the width estimation method in a vessel profile . . .	43
5-2	The four local directions we test for	45
5-3	Superimposing the identified points on the original vessel	46
5-4	Detail of original vessel	47
5-5	A part of the vasculature with the lumen defining points superimposed	47
5-6	A part of the vasculature with the wall defining points superimposed	47
5-7	A representation for our sub-pixel precision technique	49
5-8	Simulation of a vessel with a locally high curvature	50
5-9	Simulation of a vessel with a locally low curvature	50
5-10	The width of the lumen of a vessel along its length	53
5-11	Some of the original tracks after the removal of the papilla region .	53
5-12	The vessel calibre estimated on the original image	54
5-13	The total width estimate for our demo image	54
A-1	Non-maximal Suppression	62

A-2 Cleaned image	62
A-3 Dilated image	62
A-4 Eroded image	63
A-5 thinned image	63
A-6 Before removal of small tracks	64
A-7 After removing small tracks	64

Chapter 1

Introduction

Doctors are perhaps the only people that share so deep a passion with vision engineers for seeing what other people cannot normally see. For they know that the information is there, waiting to be elicited by the interpretive eye. They, like vision engineers, are faced on a daily basis with enormous loads of information, sometimes rare or unique and completely cut out of context.

However, the interpretation of information and the evaluation of its worth in a clinical environment can be of vital importance to human lives and welfare and hence cannot be left a burden on the shoulders of a passionate or talented few. The demands of health care increase day by day and the quest for positive and objective medical assessment increases along them. Doctors need to make rapid decisions on cases that are presented in front of them, often without the luxury of a literature research or extended investigations. As pointed by Szolovits [Szolovits(ed) 82],

Continued training and recertification procedures encourage the physician to keep more of the relevant information to mind, but fundamental limitations of human memory and recall, coupled with the growth of knowledge assure that most of what is known cannot be known by most individuals

Since the advent of ancient Greek theatre, solutions to pressing situations have been given by machines.

1.1 Computers in Medicine

Computers have entered the field of medicine a good time ago through the means of their successful applications in office automation and information storage. It is a common fact that no medical institution can operate without computerised archives and special programs that assist the everyday conduct and work of all the people that span health professions. But it is unfortunately this interpretation of computers as surrogate memories and tireless secretaries that has shadowed the endless possibilities of real help they could offer in the medical profession. For, what is the essential difference between a pile of paperwork and a top of the range machine, if the latter carries but unstructured bulks of information? All medium sized hospitals can speak of tons of uncorrelated facts from patients records. To the best of their expectations during the past few years these have given up their place, but not necessarily their secrets to a handful of disks. A lot of established techniques exist today for the retrieval of raw information in the form of databases but reasonable recall of information, in a form close to the needs of the clinician has yet to established. It is towards the diagnosis in the everyday medical practise that technology should turn focus, where the low number of experts could carry their expertise through well structured *information processing* systems (as opposed to *information storing*) to reach more and more people in need. Technological breakthroughs in the field of medicine are often so impressive that they make the headlines and give perhaps the false idea that cutting edge technology will sooner or later prevail over all health related problems of humanity. Of course, this is not the case in the foreseeable future. We are still at a point where a certain mass of techniques have been gathering and should soon be mature enough to constitute a well researched body of expertise, aiding the doctor in his practise and allowing him to really make the critical decisions. For there is no point in deciding when relevant information is absent or obscured.

In this state of mind we embarked on a research project that produced the document you now hold.

1.2 Overview of the Project

The aim of this project was to design and implement a non-invasive automated method for the examination and follow up of the arteriosclerotic changes in blood vessels due to hypertension, with the help of digital image processing of the retinal vessels.

Hypertension usually brings to mind the cardiovascular system of blood circulation, which has been designated as one of the most vulnerable parts of the body of modern man. However, as it happens with biological data, information is interwoven into the yet inexplicable correlations of the whole human system. It should come as no surprise that doctors would look anywhere for vital information. The human eye has long been known to be not only the mirror of soul but a clear

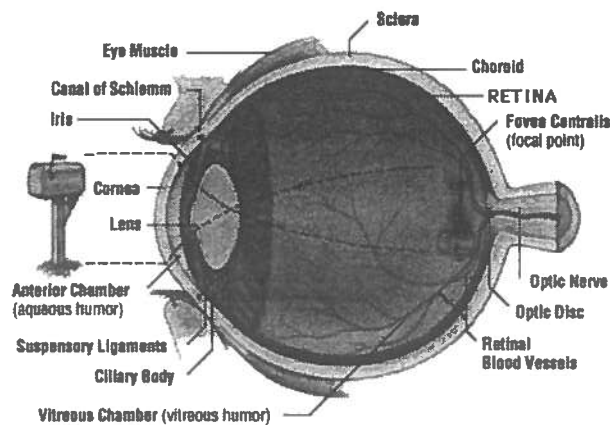


Figure 1-1: The human eye: A wealth of information

indicator of the bodily health as well. The eye presents a wealth of information and is still considered unexplored, despite the fact that techniques have developed a long time ago for the mining of that information. It is neither our purpose nor ambition to present here an extended treatise on the human eye and the possible information it could provide in the diagnosis and prognosis of systemic diseases. Nevertheless, we intend to present all relevant information that would facilitate the reader in assessing the difficulties that arise in treating biological data with

computers and help him/her clarify the terminology that governs the issues in question. The *retina*, which is the interior part at the back of our eye is the only part of the body where arteries and veins can be studied *in vivo*. Also, it has been established that changes in the retinal vessels reflect the rate of progression and severity of systemic diseases. In particular, mild hypertension causes irregularities in arterial calibre, tortuosity of the retinal arteries and changes at arteriovenous crossings. Repeated ophthalmoscopic examination provides the opportunity to observe the vascular effects of hypertension. Such ophthalmoscopic examination can be performed by the digitization of an ocular fundus picture, isolation and processing of the retinal vessels. The parameters which are of significant interest to the medical experts include:

- Apparent length of a vessel, particularly until the next bifurcation
- Calibre of the blood column
- Thickness of the wall of the vessel
- Tortuosity of the vessels

1.2.1 Objectives

We concentrated on the segmentation of the vessels from their background with a reasonable accuracy of placement. Our aim was not only to identify and isolate the vessels from the background but to do this in such a way that our later measurements would be enhanced. We felt that if we could achieve a proper and robust segmentation technique the rest of problems would benefit strongly. However, a set of methods and algorithms had to be devised and implemented in order to carry on measurements since no satisfying established method existed.

1.2.2 Achievements

An algorithm was devised and implemented for the segmentation of retinal vessels using ridge detection. The algorithm tracks the centerline of each vessel instead of the edges, a thing very useful in the later stages. It can be summarised as follows:

the image is being smoothed and is then scanned pixel per pixel in a non-maximal suppression scheme. We suppress pixels that are not local maxima or they are not contrast maxima in a wider region orthogonal to the direction of the vessel locally. This algorithm effectively gives out a skeleton of the vasculature system which is then exploited for the specific measurements. The key idea is that vessels present a ridge in cross-sectional profiles and therefore such a ridge can be detected along the vessels as a whole as if there was three-dimensional information. Once the vasculature was segmented we moved on to carrying typical measurements of interest such as width of wall and calibre of the vessel. We also calculated length and introduced a measure for tortuosity. All our measurements are calibrated by means of the papilla which is considered of standard size. To achieve this, we implemented a location technique for finding and measuring the papilla on fundus images.

1.2.3 Conclusion

The results show that ridge detection is not only a viable technique but perhaps the most recommended for our case. It leads to an ideal segmentation with respect to our interests for diagnosing and monitoring hypertension. It is a simple technique, not especially computationally expensive that could be developed in a run-time application for use in the everyday practise. These ideas are further explored in the last part of this thesis.

1.3 A Guide to the Rest of the Thesis

In Chapter 2 we provide the relevant background with respect to image processing of retinal images as well as some medical insights for retinopathy in general.

Chapter 3 describes the algorithm for the automatic location of the papilla. Two different approaches are presented and their results are illustrated on different images.

Chapter 4 contains the heart of the method developed during this project which is the tracking and subsequent automatic extraction of the vessels. Several different

approaches are presented and evaluated. The technique of ridge detection which was finally employed is presented in detail and its algorithmic levels are explained. The results of the method are illustrated on our test image.

The tracked vessels derived in the previous chapter are used in Chapter 5 for a demonstration of possible measurements that can be undertaken once the vessels are segmented. The formulae for these measurements are presented and their rationale is explained. As an example we calculate vessel calibre, wall width, length and mean tortuosity index on our test image.

Finally, Chapter 6 contains a short summary of the work undertaken for this project and an evaluation of the results. The thesis concludes with a short description of future work and extensions.

Chapter 2

Background

In recent years the eye has become the focus of medical attention for a number of systemic diseases. It is no exaggeration to say that the recognition of malignant hypertension, miliary tuberculosis, papilloedema due to increased intracranial pressure, melanoma or even diabetes may be a matter of life and death. The evaluation of a patient's condition by means of ophthalmoscopic examination requires careful visual inspection by human experts who often consider fundus angiograms. Because clinical examinations are subject to uncontrolled observer variation and bias, assessments could be made more objective using automated digital image processing. This study was motivated by research undertaken in the Medical Renal Unit of the Royal Infirmary of Edinburgh with respect to hypertension.

2.1 Medical Background

Since the eye is considered an outgrowth of the brain it is reasonable to assume that changes in the vasculature of the retina reflect changes in blood microcirculation of the brain. Indeed, as early as 1939, Keith, Wagener and Barker have proposed a classification of hypertension based on retinopathy which is still widely used [Keith *et al.* 39]. Their classification was as follows (after [Wise *et al.* 71]):

Stage I : Slight narrowing of the retinal arteries. Most patients free of symptoms.

Stage II : More marked arteriolar sclerosis with widened arterial light reflex, compression of veins at arterial crossings, general and focal arterial narrowing.

Stage III : Retinal oedema, exudates and hemorrhages with diffusely and focally narrowed arteries.

Stage IV : Papilloedema on top of symptoms of Stage III. Renal failure. Poor prognosis.

According to Wise,

“difficulties in the interpretation of vascular narrowing and arteriolar sclerosis have led many researchers to abandon the first two stages of the Keith-Wagener classification. Stages III and IV are still widely accepted as their characteristics are considered more tangible and observer independent within the medical community.” [Wise et al. 71]

As today many cases reach the physician at an earlier stage than formerly,

“arteriolar changes alone may be the presenting feature without the severer symptoms of a retinopathy such as exudates, haemorrhages or oedema.”[Stokoe 77]

Hypertension, even at an early stage, is known to manifest itself in the retina by attenuation changes, focal narrowing and occlusion of retinal vessels. Automatic quantification of these changes would be a great step not only in the direction of diagnosis of hypertension, but also towards an optimised method for monitoring the progress of subsequent treatment.

2.1.1 Vessel Calibre

The retinal vasculature can be considered a series of cylindrical pipes with respect to blood circulation [Wise et al. 71]. It is not our purpose nor ambition to explain the hydrodynamics of piping systems or blood vessels in particular. Still, we give a brief outline of the ideas that govern the pressure calculation in the blood vessels and justify our particular interest in the vessel calibre and its accurate calculation.

2.1.2 Hemodynamics

The simplest hydrodynamic system is a long straight pipe with a steady flow through it. In such a pipe the head of the pressure (Δp) is directly proportional to the length (L), the flow (F), the viscosity of the fluid (μ) and inversely to the fourth power of the radius (R). This is the law of *Poiseuille*:

$$\Delta p = \frac{8L\mu F}{\pi R^4}$$

It is therefore clear why the calibre of the vessel is important in any attempt to recognise the presence or progress of systemic hypertension. By injecting dye into retinal vessels and then studying the fluorescein angiograms it has been observed [Wise *et al.* 71] that there are two distinct types of flow in the vessels. There is "laminar" flow, where the flow can be regarded as a series of concentric laminae moving faster at the axis than away from it. The shell of fluid in contact with the wall is practically stationary. When two pipes with such flow join, dye from one of them will not mix with the other and the two streams run side by side. As the rate of flow through the pipes increases there comes a point where mixing occurs and the regular lines of laminar flow are lost. The new type of flow is no longer governed by Poiseuille's law and is termed "turbulent". The breakdown of laminar flow depends upon the diameter of the tube (D), the mean velocity of the flow (V), the density (ρ), and the viscosity (μ) of the fluid. The point of transition from laminar to turbulent flow depends upon the value of the Reynolds number (R_e). The equation from which this is calculated is:

$$R_e = \frac{VD\rho}{\mu}$$

It should be noted that the critical value of the Reynolds number for any fluid (including blood) to become turbulent is about 2000 [Wise *et al.* 71] and narrowing of the vessels alone has not been accepted as a sole reason for turbulence. Eddy formation along the walls is also mentioned as a possible culprit. However, it is established that transition to turbulence in retinal vessels always occurs close to bifurcations.

2.2 Image Processing background

2.2.1 Fundus Photography

Since the introduction of the fundus camera, many attempts have been made for the quantification of various structures in the retina. Significant improvement was brought by the development in 1961 of fluorescein fundus photography by Novotny and Alvis. The principle of this method is straightforward. Sodium fluorescein dye is injected into the circulation and photographed as it passes through the blood vessels of the eye using a blue filter in the exciting light source and a green barrier filter in front of the film [Vink 70]. For most purposes fluorescein pictures taken at a rate of one or two frames per second give satisfactory data about the filling and emptying of the vascular structures in the retina. However, the velocity of blood flow in the arterioles can only be estimated by cine-photography at a rate of 30 to 60 frames per second. An alternative approach is to film cine-angiograms directly (for a recent application of video-angiography see [Jagoe *et al.* 93] and [Jagoe *et al.* 90]). The dye is injected either intravenously, or into the anonymous artery via the right brachial artery which can result to a certain inconvenience of the patient. Furthermore, fluorescein sodium is not free of toxicity. Many patients develop a yellowish discolouration of the skin and a standard after-effect is bright yellow urine that remains fluorescent for 24 - 36 hours after the examination [Vink]. Other patients have reported a feeling of nausea and quite a few have actually vomited during the examination. Allergic shocks, usually by patients with a history of asthma, are not rare [Wise]. In short, this method may cause inconvenience to the patient and hinder the medical need for repetitive examinations in monitoring certain types of retinopathy and treatment progress. On the other hand, standard fundus photography utilises only mydriatic drops for the dilation of the pupil. The forced dilation, which is also used in the fluorescein photography, fades away soon after the examination and a few hours later the accommodation of the eye is restored completely. It is the aim of this study to indicate that for the purpose of monitoring hypertension all quantification information could be extracted from plain fundus images.

2.2.2 Previous Work on Retinal Images

The first step in performing measurements of blood vessels is to extract the vessels from their background. Estimating the location of blood vessel boundaries within images of a living vascular network “has all the problems associated with image segmentation of biological data and still remains a difficult problem” [Miles & Nuttall 93]. This is largely because such images have low signal-to-noise ratio and limited spatial resolution. Most segmentation methods for vessels are considered to be successful if a continuous boundary is found, somewhere in the vicinity of the true boundary, without much concern for boundary placement. Many algorithms have been proposed for tracing vessels and they can be generally classified in three basic categories:

- model fitting, which identifies edge points by optimising a function fitting (*e.g. a Gaussian*) on the intensity profile of a vessel.
- optimal filtering, where gray level thresholds are specified based on the distribution of pixel properties (*e.g. adaptive thresholding*)
- sequential contour tracing, where active contour models (*e.g. “snakes”*) are used to perform a global region extraction

Recently, F.P. Miles and A.V. Nuttall [Miles & Nuttall 93] studying blood flow in the cochlea of the ear, proposed an algorithm that estimates diameter and position of vessels by minimising the mismatch between the measured intensity profile and a finite set of intensity profile models, thus constructing a matched filter estimator. A similar approach is adopted by Zhou *et al* in [Zhou *et al.* 94]. Here, the centerline is tracked using an adaptive densitometric technique to improve computational performance in regions where the vessel is relatively straight. Their method calculates vessel calibre and wall location (*but not vessel wall thickness*) as well as a curvature index defined as the ratio of the pixel length of the centerline over the distance between two points. This algorithm fails at arteriovenous crossings and no information is given on its behaviour in bifurcations. Initial centerline point and direction are expected to be given by the end user.

Alternatively, T. McInerney and D. Terzopoulos [McInerney & Terzopoulos 95] have presented a topologically adaptable snakes model for angiogram segmenta-

tion. Their model benefits from reparametrization during the deformation process. Still, it suffers from the problem of initialisation common to snake models. The authors report that they had to perform the segmentation of the vascular network one branch at a time freezing the snake manually once it began to flow into a crossing branch.

2.3 Images Used in this Study

2.3.1 Image Acquisition

The images used for this study were provided in film by the Retinal Photography team of the Eye Pavilion at the Royal Infirmary of Edinburgh.

The images were captured through the use of a KOWA Pro 1 fundus camera that registers 50 degrees of the ocular fundus. The film used was black and white negative *ILFORD* FP4 Plus of 125 ASA (22 DIN). A standard green filter was used for red-free negatives. Red-free images are particularly important for size estimation in a vascular network since the photo effect influences it. The film was developed in Patterson Accuspeed Developer for 8 and a half minutes at 20 degrees centigrade.

The test image used throughout this document comes from a 32 year old white male who is not diagnosed as hypertensive and does not undergo treatment for any eye disease. The subject was brought to pupil dilation using Tropicamide 1% as mydriatic drops for 20 minutes before the photographic session which lasted less than 5 minutes. Taking into account the film processing time, the whole procedure normally takes less than 40 minutes.

A set of colour images were also provided. These were shot on *FUJI* Velvia of 50 ASA (18 DIN). They were not utilised in this study for reasons explained in the discussion paragraph.

2.3.2 Digitization

The images were digitized at the Multimedia Workshop of the Department of Computer Science of the University of Edinburgh using *ScanMaker 35t* slide scanner by Mikrotek connected to a Macintosh Performa. This scanner has a scanning area of 36 mm (+/- 1.5 mm) x 36 mm and can output up to 1828 pixels per inch (i.e. 2592 pixels total) from the original film. It can reach a resolution of up to 3656 ppi (i.e. 5184 pixels total) with interpolation. For the purposes of this study no interpolated digitization was used. The scanner allows for contrast, brightness and resolution enhancement through the driving software. Digitization takes place in one pass (three passes for colour film) and lasts about two minutes (eight minutes for colour).

The software used for the scanning was Adobe Photoshop ver 3.0 and the respective plug-in filter for this scanner.

2.3.3 Image Formats

Originally the images were saved in TIFF (Tagged Image File Format) as this is easily carried over on all different platforms. During our study images were captured in Apple Macintosh and carried over to Unix workstations (both Sun Sparc workstations and HPs) as well as MS-DOS based PCs without any problem. The black and white images take something between 2 - 6 MB of disk space depending the resolution used, while coloured ones go above 10 MB. LZW compression for TIFF images is not lossy with respect to information and reduces their size by about one third. All the manipulation of the images by the test programmes created was done on Sun workstations under Unix. To facilitate the manipulation of the images they were turned from TIFF to HIPS format using the public domain PBMUTILS utilities package enhanced by the Machine Vision Unit of the Department for HIPS support. HIPS format is essentially a raw format with a simple header structure. From that it was easy to convert the image files to plain intensity value matrices that could be directly read into Matlab for further manipulation, preserving the image dimensions.

2.3.4 Software Implementation

Throughout the project we used the Unix platform (*X Windows*) for developing software to test the proposed algorithms. Routines were implemented in Matlab 4.2 under Sun-OS4.0 and Sun-OS5.0 for Sparc 4/5 workstations. Coding for certain processes (*e.g. tracking*) had been implemented in C/C++ and was compiled under the GNU C compiler (*gcc*). Provision was taken so that the results of the implementation did not depend on the computer languages or configurations. System calls were virtually absent from the code to allow for easier portability. We wish to stress that our aim has been from the beginning not to create market-ready software but to test the validity of hypotheses and the plausibility of research application. Hence, no attempt has been made to fully optimise the code implementation with respect to time, for example.

2.3.5 Our Demonstration Image

Our demonstration image (Figure 2-1) is a cropped image from a negative B/W slide. It has been digitized at a resolution of 1828 ppi with 200% proportional magnification during digitization.




Figure 2-1: The image used for demonstration of our methods throughout this document is a 400x430 B/W fundus photograph with the papilla centred

This particular image has been selected both for its merits and its troubles. There a lot of vessels present, the papilla is well defined but there is also a good deal of unstructured noise (there is in fact a black spot in the bottom left side caused by the film negative), the background is characteristically non-uniform and some vessels seem to get absorbed inside the background intensity values. Also, the vessels that come in different lengths, widths and strengths present an interesting pattern of tortuosity, from relatively straight to sharp bending ones, as well as bifurcation.

2.4 Discussion

We decided to concentrate on the black and white images as these had an acceptable contrast with regards to vascular border definition, they were smaller and easily manipulated and they served our original purpose of extracting the vessels so as to perform the measurements of interest. On the other hand, colour images proved defective when digitized. In particular, they had an abnormally uniform value of the blue channel in the RGB digitization thus making it impossible to use them for automatic separation of veins from arteries. We believe that a different photographic technique would register the three channels more accurately and would allow for certain colour measurements. Details on such an extension to our work are given in the final chapter of this thesis. Lastly, we would like to note that the digitization step could have been rendered redundant by the use of a digital fundus camera. Such a camera was not available at RIE at the time of this study. In fact, at this time, the cost of a digital fundus camera is still prohibitive for the everyday medical practise in small to medium medical centres.

Chapter 3

Calibration

All absolute measurements undertaken in fundus images are prone to errors caused by a number of reasons. Prominent among these reasons is the refraction of light from the lens of the eye as well as the liquid that fills the area between the cornea and the retina (*vitreous humor*). Other error inducing factors are visual impairments of the eye, i.e whether the eye is emmetropic or not, as well as the sharpness of the image. Most of these errors lead to uneven magnification of the retinal structures. At the time of this study the relevant literature is still unsettled on how these errors should be calculated and accommodated. Littmann has presented a method for calculating these errors in [Littmann 82] but his methods and formulae have been criticized in [Bennet *et al.* 94] and [Baumbach *et al.* 89] and more recently by Sanchez in [Sanchez A 95]. Such errors create problems not only when monitoring the same eye over a period of time, but also in comparisons between different subjects. To minimise the impact of such errors during our study, we calibrate all our measurements according to the size of the papilla which constitutes a reference feature for our images. In order to identify the papilla we used a disk fitting procedure that is explained below.

3.1 Papilla Fitting

The papilla is known to be of relatively fixed size and it is reported as having a diameter of 1.5 ± 0.3 mm. Once the papilla is automatically located in the image, a correspondence between pixel lengths and real world measurements is established. This is reasonable since the retinal image is corrected from the lens of the ophthalmoscope for any distortion caused by the curved surface of the retina. The algorithm for the automatic location of the papilla relies on the assumptions that the papilla:

- will lie somewhere in the central part of the image
- is well approximated by a disk
- will have a radius lying between 20 and 80 pixels

These assumptions are quite acceptable in the literature. An image with the papilla centred is a basic image in ophthalmoscopic examination [Williams 84]. Additionally, in a 512×512 image a papilla with a radius between 20 and 80 pixels occupies 0.47 - 7.6 % of the image, again within reason. The test image used throughout this thesis for demonstrating the methods that were developed is a 400×430 image with a papilla of 63 pixels radius in the central region. This is the value we reached after averaging eight different diameters that were counted out explicitly for verification purposes.

For an estimation of the papilla centre we tested two different approaches. The first one fully exploits the nature and geometry of our black and white images. Since the papilla practically constitutes a hole in the posterior of the eye it was deemed reasonable that it will have the highest reflection of the light shone through the ophthalmoscope. As our greyscale image comes from a negative black and white, it should be expected that the papilla will incorporate those pixels that will have registered the lowest intensity value. Once the image was turned into an intensity value matrix and was introduced in Matlab it was fairly straightforward to track all those pixels that presented the lowest intensity value using a *min*

command. These centres were identified and their weighted centre of mass was considered to approximate the true centre of the papilla. Weights were allocated in favour of clustered points by means of a standard 3 x 3 neighbourhood weighting system to accommodate for noise. That is, for every pixel that was considered we registered its 3 x 3 neighbourhood. Now, for every pixel in that neighbourhood that had a value equal to the lowest intensity value we added one to the multiplicative factor. Note that the very pixel itself was not allowed to vote, hence giving it a multiplicative factor of zero in case it was isolated practically eliminated it in such a case. So now instead of taking

$$x_c = \sum_{i=1}^n \left(\frac{x_i}{n} \right) \text{ and } y_c = \sum_{i=1}^n \left(\frac{y_i}{n} \right) \text{ we use:}$$

$$x_c = \sum_{i=1}^n \left(\frac{x_i * k_i}{n} \right) \text{ and } y_c = \sum_{i=1}^n \left(\frac{y_i * k_i}{n} \right)$$

where n is the total number of identified pixels and k_i ranges from 0 to 7. The radius was then estimated in the same way as described below in the second method.

In the second approach a square grid of size 200 x 200 pixels around the centre of the image is considered. If we take the step of the grid being 20 pixels we can identify the centres of the 100 subsquares. This procedure yields 100 candidate points and each is tested for being the true centre of the papilla. In order to test each candidate point we fit on it a square binary mask containing a disk with the minimum radius (20 pixels) and expand the mask at reasonable steps (10 pixels) until the disk reaches the maximum radius (80 pixels). At every step we register the absolute difference:

$$\frac{\sum_{\text{area-B}} \text{out-pixels}}{\text{OUTpixels}} - \frac{\sum_{\text{area-A}} \text{in-pixels}}{\text{INpixels}}$$

where "out-pixels" are the intensity values of the pixels on the original image outside the circle, OUTpixels is the number of these pixels and respectively for area A. (Figure 3-1) The theoretical grounds for doing this is that the contrast between the inside and the outside of the disk should be maximum when the mask overlaps the true centre of the papilla. Convolving this mask at the preselected points yields an estimate of the centre of the papilla. Once the point with the highest response with respect to contrast maximization is identified, we perform a more accurate estimation of the radius by refining the increment step of the tentative radii with the disk being fitted on the designated point. For this method




Figure 3-1: Contrast enhancement between two areas that the mask looks at




Figure 3-2: Papilla location with the first method proposed

to be effective, provision is taken in the construction of the mask so that the number of pixels inside the disk is roughly the same, and hence comparable, with those that lie outside.

3.2 Experimental Results

For our test image, both methods have been applied and the results can be viewed in Figures 3-2 and 3-3 respectively.

Here we have used the results of the locating functions to overlay a circle, of the estimated centre and radius on the original image. The first method takes a few seconds to estimate the location of the centre and about 8 minutes to refine the radius on a SPARC-5 workstation. It reported a radius of 66 pixels as an




Figure 3–3: Papilla location with the second method proposed




Figure 3–4: Papilla location in a different image with the second method

estimation. On the other hand, the second method takes a bit more than 5 minutes to report an estimation of the centre and since it follows the same routine for estimating the radius, another 8 minutes for that. The estimate it gave for our demonstration image was 64 pixels, less than 2% deviation from the radius counted explicitly. Additionally, the second method has been applied on an image with different geometry (517 x 613) which fulfils the fitting assumptions mentioned earlier. The results of this experiment can be viewed in Figure 3–4.

3.3 Discussion

Both methods for the automatic location of the centre of the papilla worked well. The first method relies heavily on *a priori* knowledge of the nature of the image and despite the weights used it is expected to suffer in extreme noise conditions or in the presence of artifacts in the image due to the medical condition of the

subject. However, it allows for an extremely fast implementation and behaves robustly in “standard” noise conditions.

The second method presents a trade-off between speed and accuracy of estimation for the centre of the papilla. The more refined the grid, the more candidate points arise and hence the longer it takes for the location of the best approximation. However, this study is concerned with the development of a method to be used in off-line processing of fundus images in order to assist in the diagnosis of long term hypertension and therefore time (in the scale of seconds or a few minutes) is not of critical importance. Essentially, this second method is a sort of intelligent global optimisation using knowledge of the image as a shortcut. Instead of that, any other global optimisation method could be used. However, some of them, like “blind” template matching, would present problems of computational efficiency. By “blind” template matching we mean convolving the whole image with all possible disk sizes. Knowledge of the image is crucial in the speed of the method.

What is of particular interest is the fact that with both methods the accuracy in the estimation of the size of the radius is not influenced greatly by the accuracy in the estimation of the centre of the papilla. This is due to the very nature of the locating algorithm: since the algorithm looks for the maximum of contrast even if one goes a few pixels off-centre, the correct radius is reported. If for some reason we are interested in greater accuracy, for example we want to remove the papilla before conducting our measurements, we could also refine the centre estimation by repeating the process of testing different centres, confined this time in a smaller area derived from the first estimation. Also, one should notice that the vessel structure inside the papilla does not mislead the algorithm because their overall intensity contribution is always small compared to the darker part of the papilla.

Chapter 4

Extraction of retinal vessels

In order to measure the retinal vessels we need to segment the vessels from the rest of the image. This proved to be more difficult than it sounds, since in the case of fundus photography it is uncommon for the background and the vessels to have distinctly different ranges of intensity values. Furthermore, due to the nature of our study no chemical contrast enhancer has been used, as in typical fluorescein angiograms, which would make the vessels stand out clearly from the background. A variety of approaches have been considered, including the following:

- Thresholding the image with the use of a histogram
- Traditional edge-detection (*e.g. Canny*)
- Ridge point detection across the width of the vessels.

The first two proved inadequate for our problem for reasons that will be explained below. Ridge detection, on the other hand, gave very good results. This technique and its algorithmic implementation constitute the main part of this chapter.




Figure 4-1: A histogram of our image

4.1 Initial Approaches

4.1.1 Thresholding

By looking at the histogram of our test image (Figure 4-1) it should come as no surprise that standard thresholding techniques do not hold for our case.

The translucent nature of the vessels' walls as well as the dispersion of the light coming from the ophthalmoscope on the non-planar surface of the retina make a good part of the vessels being digitized present intensity levels close or identical to those of the background. Also the background is not of uniform intensity, something that makes the segmentation of the image in foreground and background practically impossible. All attempts on identifying a suitable threshold value have led to unacceptable results similar to those that can be seen in Figure 4-2. Hence the thresholding approach was abandoned and focus was switched to traditional edge detection.

4.1.2 Edge Detecting operators

The most prominent among edge detectors is undoubtedly the one presented by Canny [Canny 86]. In short, there are four main phases in the Canny edge detector:

- Convoluting the image with a Gaussian to reduce noise and remove small features




Figure 4-2: An attempt to threshold

- Gradient estimation
- Non-maximal suppression, i.e. elimination of points that are not local maxima with respect to gradient magnitude
- Hysteresis tracking, which utilises a double threshold to further limit non-edge points

There are a number of problems that arise with the use of the standard Canny edge detector (SCED) in our case. First of all, the SCED detects *edges* and therefore gives the outside walls of the vessels. This, however, could create serious problems in exact measurements of vessel properties such as wall thickness and vessel calibre. The third phase of the SCED, tracking, entails in practice the problems of *threshold setting* outlined earlier in this section: Since tracking relies on setting two threshold values for upper and lower threshold we finally end up with fragmented vessels as can be seen in Figure 4-3.

Notice that not only quite a few of the vessels appear broken but, also, the SCED has detected edges inside the vessels. This can be accounted for by the difference between the intensity values of the vessel walls and the blood stream that runs through them. This is an inherent characteristic of fundus images that was utilised in the ridge detection approach that was finally employed. Interestingly, the SCED does not merge vessels that are very close to one another, despite the use of Gaussian smoothing. Lastly, it should be noted that the Canny edge detector implementation used in our example is one that does not fail at junctions by false




Figure 4-3: Results of Canny edge detector

suppression of points on the weaker edge that reaches the junction. This problem has been treated in the manner explained in [Li *et al.* 89].

4.2 Ridge Detection

Ideally we would wish to obtain a single-pixel wide line running through the centre of each vessel. This is the basis for tracking the vessels and simplify our measurements. However, a cross-sectional intensity profile of a typical vessel (Figure 4-4) indicates that vessels present two peaks separated by a valley. This is due to the variation in intensity caused by the blood stream inside the vessel and the different materials of the cell walls.

To overcome this a Gaussian smoothing is introduced (Figure 4-5). Once the image is smoothed a directional map is being built by scanning the image and registering the direction of the gradient at each point.

The image is scanned again and ridge points are identified by suppressing pixels which do not satisfy the following criteria:

- Directional consistency with their neighbouring pixels
- Intensity maximization over their neighbouring pixels in the direction orthogonal to the local direction of the vessel

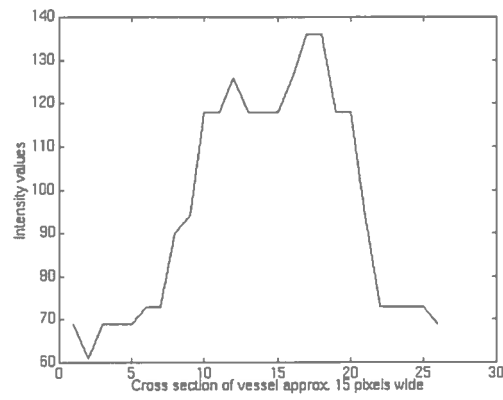


Figure 4-4: A vessel profile before smoothing

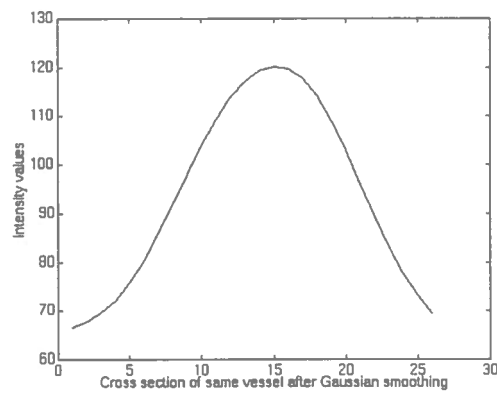


Figure 4-5: Vessel of figure 4-4 after smoothing

- Contrast maximization in the direction orthogonal to the local direction of the vessel

During the ridge point identification, the image is binarized and after that it undergoes a filtering process to clean some noise by removing individual pixels. Dilation and erosion cater for fragmented vessels. At this point thinning filters prepare the image for tracking. The vessels are then tracked and a list containing pixel co-ordinates is produced. A second filtering of vessels under a certain length practically eliminates noise completely and concludes the segmentation.

This method has proven to be quite efficient and robust and was the one finally employed. Its different algorithmic stages are illustrated in Figure 4-6

4.2.1 Smoothing

A cross-sectional profile of intensity along the length of a vessel indicates that it would be quite difficult to locate points on the central line due to noise and the variation in intensity the blood stream presents even within the same vessel (Figure 4-7). To overcome this problem we introduce Gaussian smoothing.

As has been observed by [[Fishler & Wolf 83]] Gaussian smoothing may well play the role of a distance transform when it comes to the detection of linear structures. At this point, not only are we not concerned with the possible blurring of edges but indeed we would like to eliminate or blur all structure or texture inside the vessels. In picking an appropriate width for our filter we need only take into account the possible diameters of our vessels. These are not expected to exceed 20 pixels and therefore, a Gaussian smoothing with Gaussian standard deviation equal to 19 was thought appropriate. The results of this on the vessel whose profile was illustrated in Figure 4-7 can be viewed in Figure 4-8 where clearly the valley has been turned into a ridge and respective ridge points can now be detected.

Gaussian smoothing works very well across the entire image irrespective of the direction of the vessels (see Figures 4-9 and 4-10) and does not fail at vessel bifurcations. (Figure 4-11)

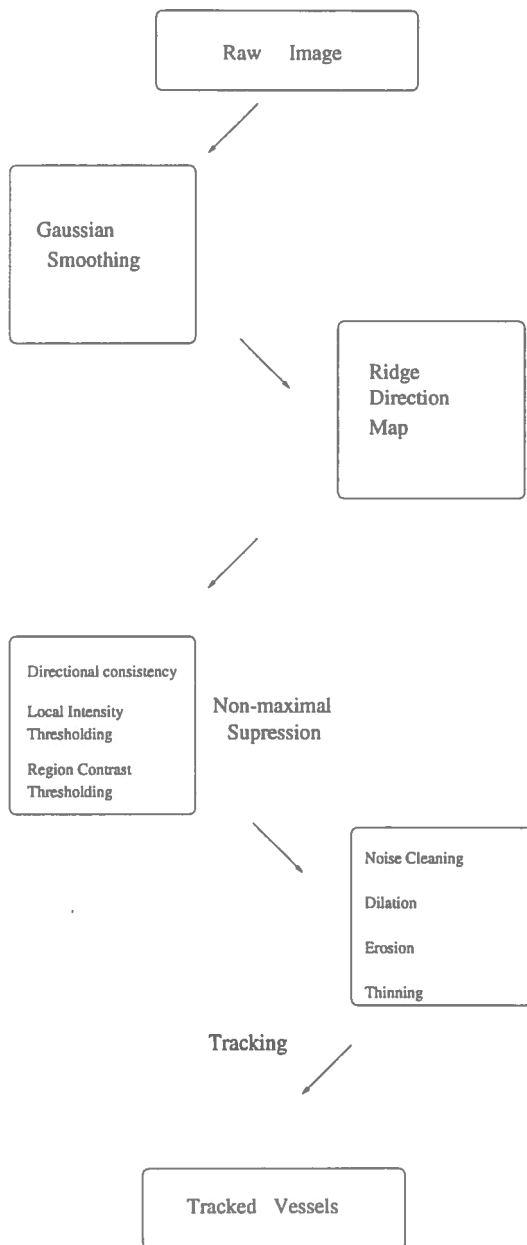


Figure 4-6: The steps of our method for vessel extraction




Figure 4-7: A typical vessel, forming clearly a valley of intensity values




Figure 4-8: Vessel from Figure 4-7 smoothed. A ridge has now taken the place of the valley




Figure 4-9: A horizontal vessel before smoothing




Figure 4-10: The vessel from Figure 4-9, after smoothing




Figure 4-11: The Gaussian is not confused at vessel bifurcations

4.2.2 Non-maximal suppression

In order to isolate ridge points we need to scan through the image and suppress all points which do not qualify as such. The first criterion we should employ is directional consistency. This means that for a point to be called a ridge point it must have the same ridge direction with its immediate neighbours in the direction that the vessel runs. We can control this by building and consulting a ridge direction map of the image.

Since the direction of a vessel varies smoothly, we build our direction map by convolving our smoothed image with four different ridge detection masks, one for each main direction (0, 45, 90 and 135 degrees) and registering the response of each pixel. The mask with the highest response corresponds to the direction of the gradient at this particular pixel. The masks we used are shown below :

-1	2	-1
-1	2	-1
-1	2	-1

-1	-1	2
-1	2	-1
2	-1	-1

-1	-1	-1
2	2	2
-1	-1	-1

2	-1	-1
-1	2	-1
-1	-1	2

Note how the coefficients of the masks add up to zero, minimising the directional response of near homogeneous regions. The directional map for our test image can be viewed in Figure 4-12. Colour codes correspond to the four different directions. We built our directional map with the value 63 corresponding to 0 degrees, 127 for 45 degrees, 191 for 90 degrees and 255 for 135 degrees. The directional map presents a number of problems. Noise on the ridge causes misclassification of directions. This leads to rejection of valid ridge points at later stages




Figure 4-12: Ridge direction map

where directional consistency is tested for neighbouring points and therefore, to fragmentation of vessels. Also, we noticed side bands appearing bilaterally to each vessel. Since we are using ridge detectors this can be accounted for by the behaviour of our masks on the slopes of the ridges. As these areas have a response close to zero for all directions, only the order in which the responses are checked can influence the colour coding. Still, the vessels stand out clearly in this map and the information obtained can be utilised for our non-maxima suppression.

4.2.3 Local and Contrast Thresholding

To further suppress false local maxima we need to look at the total intensity variation not only in the immediate neighbour of each pixel, but in the broader region as well. This is because when detecting ridges we do not want to be misled by local maxima due to noise. We do this by ensuring that, in a direction orthogonal to the one registered in the direction map, our candidate ridge point has a higher intensity value than its immediate left and right points.

Furthermore, we check that a point to be accepted must have an intensity value higher by a set value (equal to 4 in our experiments) than points located at a set distance to the left and to the right, again in a direction orthogonal to the vessel. This set distance depends on the diameter of the vessels and during our study it was determined to be 6 pixels. Points that do not qualify with respect to these two test are eliminated. The results of this filtering can be viewed in Figure A-1

4.13

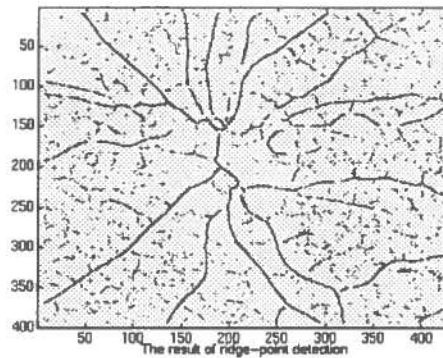


Figure 4-13: Non-maximal Suppression

4.2.4 Preparation for Tracking

Despite the tests that has been done so far, many false maxima are still present in our image mainly due to noise and small vessels. Additionally, since we are looking at ridge points, noise on the ridge of a vessel can cause false registration of ridge direction and lead to suppression of true ridge points and hence fragmentation of vessels. Lastly, the centerline that has been detected for each vessel is not necessarily one pixel wide (which we would like for the tracking algorithm to work correctly). We now have to take our image through a set of relevant operators to accommodate for these problems.

We first convolve the image with a cleaning filter which removes isolated ridge points. Since at this stage we have a binary image this mask could be a 3x3 neighbourhood mask that sets to zero all those pixels which have no other active pixel in their neighbourhood. Next comes a standard dilation convolution filter which can again be implemented as a 3x3 mask that sets a pixel to 1 if at least one other pixel in its neighbourhood is active. Last in this set of filters comes an erosion filter. Once again this can be implemented as a 3x3 mask that sets a pixel to zero if at least one of its neighbours is zero. This way fragmentation is kept to a minimum. Our image is now ready for tracking save for the fact that the tracker works best with thinned images. Thinning is in effect an iterative convolution with different sets of masks which remove the middle pixel if they get a perfect match. Many times these masks contain wild-cards, meaning that they would accept fine a zero or a one in the place of the star. For our image we used

two HIPS commands, `lthin2` and `lthin1` which convolve the image with masks like those we can see below:

0	0	*
0	*	1
*	1	1

0	0	0
1	*	0
0	1	1

0	0	*
0	*	1
*	1	1

0	0	0
1	*	0
0	1	1

1	1	1
1	*	1
0	0	0

0	0	1
0	*	1
0	0	1

We first applied `lthin2` which performed 7 iterations and then `lthin1` which performed only one.

4.3 Tracking

Now that the image is thinned we can begin tracking the centerlines of the vessels. To do this in our study we employed a classic edge tracker proposed by Beattie[Beattie 82] optimised for the treatment of bifurcations. This algorithm scans the whole image and utilises a set of pixel co-ordinate stacks for tracking. An active pixel in our image can fall in one of four categories. It can either be




Figure 4–14: Our image thinned and ready for tracking

- a starting point
- a midpoint
- an endpoint
- a start-point for a new track bifurcating from the one being followed

If the pixel scanned is a starting point it is marked as visited and its track is followed. The same happens with midpoints. When a start-point of a new track is met it is pushed onto a stack for later treatment and the algorithm follows the original track. Once the endpoint is reached any start-points that have been met during tracking are popped from the stack and their track is treated recursively. Marking of visited points prevents infinite loops. When no more points can be popped the algorithm continues the scanning of the image. The implementation we used was optimised for short-length sharp changes of direction. These are not confused for curves of the major tracks but are registered as bifurcations to small arterioles or venules. This facilitates our interest in counting mean length to the next bifurcation and prevents errors in the calculation of a vessel's mean tortuosity index.

This particular tracker produces a file in tracked edge format, that is after a simple header it has a list of entries of the form:




Figure 4-15: The final tracks are superimposed on the original grey level image.

```
n1 xxx yyy zzz
```

where $n1$ can be either 0,1,2 or 3, xxx and yyy are the co-ordinates of the tracked pixel and zzz is it's intensity value. The first number, $n1$, denotes a starting point of a new track if it is equal to 1 and a midpoint or endpoint if it is equal to 2. The values of 0 and 3 are reserved to denote the start and end of the tracked edge information respectively. The numbers that appear next to 0 in the place of co-ordinates are the size of the image while those that appear next to 3 random and can be ignored.

After the track list is complete we pass it through a special filter that eliminates tracks of length smaller than a set number of pixels. In our study it was deemed that tracks with total length under 10 pixels carry too little significant information. You can view the results of our tracking superimposed on the original gray level image in Figure 4-15. Here we have removed all tracks with lengths under 10 pixels, but still spurious tracks remain. Also, it is apparent that we have a problem of fragmentation. Notice, however that the tracking algorithm detects well the centerline of the vessel as can be seen in Figure 4-16 that shows a magnified part of Figure 4-15. This is a strong indication that the ridge detection technique is appropriate for our problem. Also, notice that the impression for the extent of fragmentation given by the printed image in Figure 4-15 is false. Due to the fact




Figure 4-16: A magnified part of three tracked vessels




Figure 4-17: Tracking results in a different image




Figure 4-18: Detail of an arteriovenous crossing

that our tracks are one-pixel wide it is expected that halftone printers and raster screens will not depict them completely in low resolutions. However, recombination and retracking of the long vessels is necessary, as we point out in the discussion that follows. An image of different geometry is presented in Figure 4-17 for comparison. Notice how our method does not fail at arteriovenous crossings (Figure 4-18).

4.4 Discussion

During our study the method described above successfully extracted 80% or more of the vessels present in a fundus image compared to those that can be seen with the naked eye. If we restrict ourselves to large vessels the percentage of success is well above 90%. However, a number of problems were identified:

- merging of nearby vessels with strong intensity variation
- fragmentation of tracked vessels

One particular case where merging occurs in our test image is shown in Figure 4-19. Here the right vessel has registered a very high intensity compared to the one immediately to the left. As a result any attempt of Gaussian smoothing appropriate for ridge detection will lead to the centerline of the left vessel being suppressed.




Figure 4–19: A fraction of our test image

In Figure 4–20 we can see the original intensity values of the region and three different attempts to smooth it. The one used in the results of Figure 4–15 is the one with width equal to 19. Clearly the left vessel does not have a ridge with the assumptions set forward in our study since the candidate ridge point is not intensity maximal compared to its immediate neighbours in a direction orthogonal to the direction of the vessel. If we lower the Gaussian width, we still do not get a detectable ridge point since now our candidate point does not present a contrast maximum in a direction orthogonal to the direction of the vessel. Decreasing the width further leads to noisy results which completely fall out of our ridge detection technique as can be viewed in the $\sigma = 5$ smoothing plot. It should be noted that the intensity level of this vessel *per se* does not exclude it from being identified, as the one on the far left of Figure 4–19 is, while having a similar intensity profile.

The fragmentation of vessels proved quite hard to avoid, the reason being that the requirement of directional consistency on a pixel-thick line is prone to noise. It is interesting that the set value which determined the distance over which a candidate point is checked for being a regional contrast maximum did not influence greatly the fragmentation. The optimisation of our tracking algorithm has the tracker stop at sharp changes of direction and register a new track for these. This was thought appropriate for the later requirements of calculating mean length until the next bifurcation and mean tortuosity index. Of course, the tracker cannot unfragment vessels which have been falsely fragmented in previous steps. Still, we

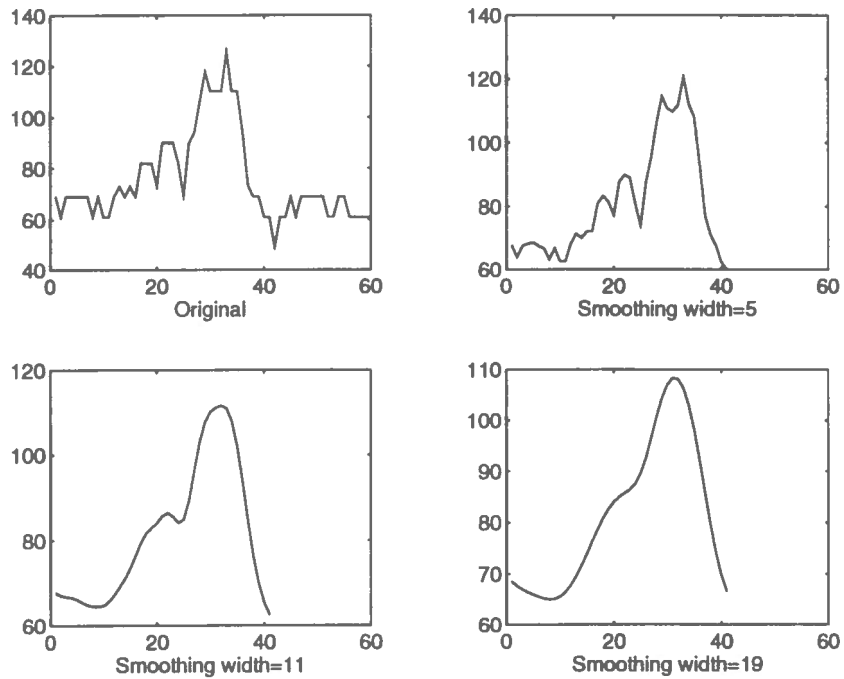


Figure 4–20: Different levels of smoothing on a vessel

think that an intelligent linking procedure should be developed for joining fragments of long vessels and identify bifurcations. This would provide an overall view of the vessels instead of the local view adopted in this study and could provide the ability to perform global measurements of length and tortuosity. Several attempts performed in the direction of track linking utilised algorithms that rely on calculating mean positions in clusters of points and interpolate intermediate points to perform joining. Results of such a technique are shown in Figure 4–21 superimposed on the original image and are of course rejected in our case as they distort the accuracy of subsequent measurements.




Figure 4-21: Linking by interpolation

Chapter 5

Measurements of Retinal Structures

In this chapter we use the vessel extraction derived earlier in order to make estimates for sizes of retinal structures. We are mainly interested in the length of vessels, the width of the vessel lumen (hereafter referred to as “vessel calibre”) and the width of the wall of the vessel. We also introduce a measure of tortuosity for the vessels and calculate all of the above for our demonstration image.

5.1 Arteriolar length

We estimate the length of a vessel by consulting the tracked file and counting the pixels for each individual track. This approach is under the assumption that the length of the centerline is a good approximation of the length of the vessel and that our tracker stops at bifurcations and gives a new track for every sub-vessel. We note that of medical interest in our case is not the *total* length of a vessel but the length until the next bifurcation. This allows for monitoring hypertensive retinopathy which makes vessels bifurcate earlier than they do in normotensive persons [Stokoe 77].

The length of the tracks is reported as the tracks are met in the track file. A typical output of our length counting program looks like this:

```
Track # 1 length 13
Track # 2 length 86
Track # 3 length 107
```




Figure 5-1: Application of the width estimation method in a vessel profile

- Track # 4 length 34
- Track # 5 length 87
- Track # 6 length 21
- Track # 7 length 58
- Track # 8 length 18
- Track # 9 length 32
-

Currently we have implemented a program that overlays the tracks on the original greyscale image reporting at the same time their length on the screen. Other configurations that would enhance the visualisation of such results, such as colour-coding, are possible. This is a topic that will recur in the last chapter, where the extensions to this work will be discussed.

5.2 Estimation of Vessel Calibre and Wall

Apart from the significance of the vessel calibre changes that lead to (and depend on) pressure fluctuations, it is apparent that the width of the wall is an interesting parameter in monitoring atherosclerotic changes that result from hypertension. Since at this point we have identified the pixel co-ordinates of the points that constitute the centerline at each vessel, we can use these to determine automatically the width. We do this by first sampling the original image in a direction orthogonal to the local direction of the vessel. Since the track is one-pixel wide,

this direction can be derived by comparing the relative position of the scanned point with respect to its neighbours. We also sample the same points from the smoothed image. Sampling goes out a set number of pixels to the “left” and to the “right” of the scanned pixel in a direction orthogonal to the local direction of the vessel. This set number during our experiments has been determined to be six pixels in the horizontal or vertical direction and five pixels in the two diagonal directions (45 and 135 degrees) to accommodate for the rectangular shape of the pixels. We subdivide our samples to “left” and “right” pixels and identify where the difference between the original and the smoothed image becomes minimum on each side. The rationale for doing this is that the points less affected by the Gaussian smoothing (in the region sampled) will be the ones that lie closer to the border of a blood vessel. Additionally we identify the two maxima to the left and to the right of each point registered in the tracked file. These are taken to correspond to the beginning of the lumen of the vessel in either side. Hence, calculation of the width of the wall of the vessel has now been reduced to calculating the difference between the pixel positions of the points that mark the walls of the vessel and those that mark the beginning of the lumen. In Figure 5-1 we can see the application of our method on a cross-sectional profile of a vessel that appears in our test image. It may seem that the points identified, particularly in the outer walls, as border points are not the ones where the smoothed image is close to the original, but this is because our polts have been done with interpolated points, while our values are discrete. For demonstrative reasons we have overlaid the detected pixels on the original greyscale image corresponding to this particular vessel (Figure 5-3). Figure 5-4 is a further magnification of this part of the image. We also include two greyscale images of a wider area. Figure 5-5 depicts the result of overlaying pixels identified as lumen borders inside the vessel. Figure 5-6 depicts the same area but this time with the pixels that are estimated to correspond with the wall of the vessels. Full scale overlaid images are to be found in Appendix A. The algorithm for estimating the vessel calibre and wall thickness is illustrated below. Note that w is the size of the set value of pixels we sample. The directions we refer to can be viewed in Figure 5-2 :

```
READ track_file  
  for every track  
    while this track is not exhausted
```




Figure 5-2: The four local directions we test for

move to the next point

register two points:

current point $P(x,y)$

next point $N(n,m)$

if $x=n$ then

direction to look at is 90 deg.

sample original image orthogonal to direction

sample smoothed image orthogonal to direction

*end * if **

if $y=m$ then

direction to look at is 0 deg.

sample original...

sample smoothed...

*end *if **

if $(x < n \ \& \ y < m)$ OR $(x > n \ \& \ y > m)$ then

direction to look at is 45 deg.

*sample original... *one sample less on each side **

*sample smoothed... * one sample less on each side **

*end * if **

if $(x < n \ \& \ y > m)$ OR $(x > n \ \& \ y < m)$ then

direction to look at is 135 deg.




Figure 5-3: Superimposing the identified points on the original vessel

```

    sample original... *one sample less on each side *
    sample smoothed... *one sample less on each side *
end *if *
    build subvectors of length  $\frac{w-1}{2}$ 
*Store the 'left' and the 'right' pixels from the one
scanned. Note that the pixel scanned is not included in these
subvectors. Also, w must be odd.*
    find the maximum of each subvector and store the relevant pixels
    calculate the respective differences
    built the respective difference subvectors
    find the minimum of each subvector and store the relevant pixels
end
END

```

Once the co-ordinates of those pixels are stored, the width of the vessel at every point that has been recorded by the tracker can be estimated using, for example, the Euclidean distance between the two external pixels that have been calculated for this particular point. This way we can have a detailed width variation of each particular vessel along its length as well as an average width of all the vessels tracked. Also, we can easily estimate the wall width as the distance between the pixel where the sample from the original has a maximum and the respective pixel that was found to represent the wall on that side.




Figure 5-4: Detail of original vessel




Figure 5-5: A part of the vasculature with the lumen defining points superimposed




Figure 5-6: A part of the vasculature with the wall defining points superimposed

5.2.1 Sub-pixel Precision

It is possible to attain sub-pixel precision in calculating both the total width and the wall width of a vessel. In order to achieve this instead of the absolute point-to-point difference between the samples of the smoothed and the original image we consider the standard difference and we look for zero-crossings in this difference. A realistic example from our demonstration image might help to clarify this estimation.

One cross-sectional profile of a vessel in our demonstration image has the following intensity values:

82	82	82	82	82	90	94
94	94	94	90	82	82	

Their respective smoothed values are:

80.3714	81.5666	83.1959	85.2713	87.6213	89.8734	91.5490
92.2259	91.6854	89.9774	87.3967	84.3881	81.4207	

which means that their difference vector is:

1.6286	0.4334	-1.1959	-3.2713	-5.6213	0.1266	2.4510
1.7741	2.3146	4.0226	2.6033	-2.3881	0.5793	

Now we are looking for zero-crossings (i.e. pixels where the difference changes sign) and we favour the furthest zero-crossings on each side. According to our rule of sampling the seventh element of our vector corresponds to the ridge point registered in the track file so we look to the “left” and to the “right” of it for the pixels that will determine the vessel calibre and the total width of the vessel. Incidentally, this is the profile of the vessel depicted in Figure 5-1 so it is a vertical vessel and the terms *left* and *right* here have their usual meaning on the horizontal field of view. However, this is just to facilitate our example and does not decrease the generality of our method. The furthest zero-crossing to the left appears between the second and the third pixel which have registered difference of 0.4334

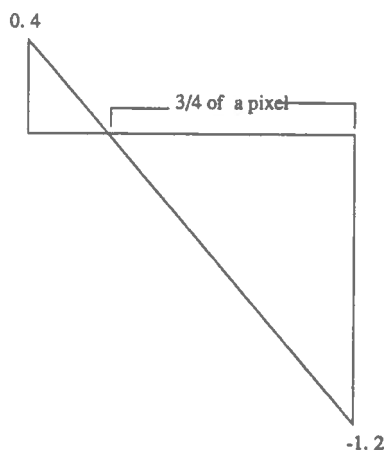


Figure 5-7: A representation for our sub-pixel precision technique

and -1.2454 respectively. Since $abs(-1.2454) + abs(0.4334) \cong 1.6$ we can say that the best approximation for the border of the vessel lies at the first quarter of the second pixel. Hence, instead of considering the wall of the vessel to start 4 pixels to the left of the ridge point, we take it to start at 4.734 pixels. A different way to locate the zero-crossing with sub-pixel precision would be the following: Suppose that the zero-crossing occurs between pixel x with value V_x and $x+1$ with value V_{x+1} . Then the zero-crossing is at $x + \frac{\|V_x\|}{\|V_x\| + \|V_{x+1}\|}$. The other side is treated in a similar manner. Of course, such an approach has no positive impact on our ability to visualise the border on a computer screen but it helps us in the accuracy of our measurements. Figure 5-7 presents a visualisation of our rationale for preserving the analogy of each pixel's "strength" with respect to the exact border.

5.3 Tortuosity

Zhou *et al* used as measure of tortuosity the ratio of centerline length between two points in question to the distance between these points [Zhou *et al.* 94]. They do not mention what metric was used for the derivation of the distance, but we may safely assume that it was either Euclidean or the Manhattan metric. In either case, such a measure is not robust as it is not able to discriminate clearly between the cases of tortuosity depicted in Figures 5-8 and 5-9. In our opinion a measure of tortuosity for retinal vessels should register a higher tortuosity index for the curve




Figure 5-8: Simulation of a vessel with a locally high curvature



Figure 5-9: Simulation of a vessel with a locally low curvature

in Figure 5-8 as its medical impact is higher. A possible measure of tortuosity in our case could be produced by averaging the change of angle calculated at reasonable discrete steps along the length of the vessel. The change of angle is independent of scale and does not cancel out along the length of the vessel. We therefore introduce as mean tortuosity index of a track, T_t , a number derived in the following manner: for each pixel indicated in the track list, P , we consider two more pixels, $P-s$ and $P+s$ that lie a set number of pixels ahead and after P , respectively. We then form the vectors $(P-s, P)$ and $(P, P+s)$ and we normalise them by dividing each with its norm. Lastly, we form their dot product and take the inverse cosine of this product. If we average these angles over the number of points used along the vessel track we get the mean tortuosity index of the respective track. This index is not reported when very few points are sampled (e.g ten or less). In mathematical notation, the ideas expressed above can be formulated as:

$$\frac{1}{(t_length - 2 * step)} * \sum_{n=step}^{(t_length - step)} \arccos(UV(P_{n-step}, P_n) \bullet UV(P_n, P_{n+step}))$$

where t_length stands for the length of the particular track and UV means unary vector. A detailed presentation of our algorithm is the following:

```

READ track_file
  for every track

```

```

repeat
  move at point P, "step" pixels along the track
  if P + step lies outside the track
    break
  increase counter
  register points for calculation:
    P := (xp, yp)
    P-step := (x1, y1)
    P+step := (x2, y2)
  calculate unary vectors:
    (V1, V2) := [  $\frac{(x_p - x_1)}{\sqrt{(x_p - x_1)^2 + (y_p - y_1)^2}}$ ,  $\frac{(y_p - y_1)}{\sqrt{(x_p - x_1)^2 + (y_p - y_1)^2}}$  ]
    (U1, U2) := [  $\frac{(x_2 - x_p)}{\sqrt{(x_2 - x_p)^2 + (y_2 - y_p)^2}}$ ,  $\frac{(y_2 - y_p)}{\sqrt{(x_2 - x_p)^2 + (y_2 - y_p)^2}}$  ]
  calculate the arccos(V1 * U1 + V2 * U2)
  accumulate the results in a variable T
calculate Ti of this track by dividing T with the counter
end
END

```

5.4 Discussion

With this kind of tortuosity measure we can also have an idea of variation of tortuosity along a vessel. This could be used in the detection of abrupt changes that could correspond to stenoses or the presence of retinal lesions. The format of the results we get is:

```

step =
  15

```

```

Track # 1 Average tortuosity N/A Very few points tested 1
Track # 2 Average tortuosity 0.1235 # points tested 56
Track # 3 Average tortuosity 0.2432 # points tested 77
Track # 4 Average tortuosity N/A Very few points tested 4

```

Track # 5 Average tortuosity 0.0923 # points tested 57
Track # 6 Average tortuosity N/A Very few points tested 1
Track # 7 Average tortuosity 0.1199 # points tested 179
Track # 8 Average tortuosity N/A Very few points tested 1
Track # 9 Average tortuosity 0.1391 # points tested 127
Track # 10 Average tortuosity N/A Very few points tested 1
Track # 11 Average tortuosity N/A Very few points tested 1
Track # 12 Average tortuosity N/A Very few points tested 1
Track # 13 Average tortuosity N/A Very few points tested 1
Track # 14 Average tortuosity N/A Very few points tested 1
Track # 15 Average tortuosity N/A Very few points tested 1
Track # 16 Average tortuosity 0.1755 # points tested 102
Track # 17 Average tortuosity 0.0695 # points tested 57
Track # 18 Average tortuosity N/A Very few points tested 2
Track # 19 Average tortuosity 0.0908 # points tested 30
Track # 20 Average tortuosity 0.0423 # points tested 27
Track # 21 Average tortuosity 0.1513 # points tested 120
Track # 22 Average tortuosity 0.2910 # points tested 151
Track # 23 Average tortuosity 0.1209 # points tested 22
Track # 24 Average tortuosity 0.2372 # points tested 33
Track # 25 Average tortuosity 0.1170 # points tested 21
Track # 26 Average tortuosity N/A Very few points tested 1
Track # 27 Average tortuosity 0.1393 # points tested 24
Track # 28 Average tortuosity N/A Very few points tested 5
Track # 29 Average tortuosity N/A Very few points tested 10
Track # 30 Average tortuosity 0.3418 # points tested 92
Track # 31 Average tortuosity N/A Very few points tested 7
Track # 32 Average tortuosity 0.3522 # points tested 38
Track # 33 Average tortuosity N/A Very few points tested 1
Track # 34 Average tortuosity N/A Very few points tested 1
Track # 35 Average tortuosity N/A Very few points tested 1




Figure 5-10: The width of the lumen of a vessel along its length




Figure 5-11: Some of the original tracks after the removal of the papilla region

Compare the results with the image in Figure 5-11. Here we have numbered the tracks to indicate correspondence with results and we have removed the papilla region because vessels twist and bend a lot on the verge of coming out of the optic disk, hence we would have accepted false indications of tortuosity as true. Tracks 30 and 32 have a markedly higher index of tortuosity which is reasonable given that their curves are more sharp than, say, track 16.

We can see that our aim to find a measure of tortuosity that will reflect the extent and abruptness of twist and turn of the vessels has been met successfully.

On the other hand, the width estimating technique has not worked perfectly. There is a great fluctuation of estimated width along the length of a vessel. We




Figure 5-12: The vessel calibre estimated on the original image




Figure 5-13: The total width estimate for our demo image

calculate this by considering the Euclidean distance of the points that are found furthest on either side of the track point in a direction orthogonal to the vessel. There are two sizes that we can calculate this way. One is the width of the vessel and the other is the width of the lumen, in other words the vessel calibre. In Figure 5-10 we can see the trembling effect created by the diversions in the estimation of the width. We close this chapter by introducing the results of our estimates on the whole image both for the vessel calibre (Figure 5-12) and the total width (Figure 5-13).

Chapter 6

Conclusions

6.1 Summary of goals

The main objective of this project was to devise and implement an algorithm for the automatic extraction of retinal vessels in fundus images. This could be a first step in a more general processing session where measurements of retinal characteristics and comparisons to known values could be made. It was clear from the beginning that later stages influence the method selected to perform the segmentation as there is no guarantee that a method which segments the vessels from the background however sophisticated and mathematically sound (e.g. Topologically adaptive snakes [McInerney & Terzopoulos 95]), will allow for further processing of the image features in the desired direction. Once the vessels were accessible we were required to perform a number of measurements with respect to sizes of different structures on the retina. As indicative of the problems that arise when dealing with biological data we were asked to calculate length, width of vessels and wall thickness. Furthermore, we were required to introduce a tortuosity index measure that would give an idea about the correlation of systemic hypertension with attenuation changes in retinal vessels.

6.2 Summary of the work completed

During this project we developed and tested an algorithm for the extraction of retinal vessels based on the technique of ridge detection. Our technique exploits the fact that vessels in fundus images have a ridge appearance with respect to intensity values. These ridge points were detected by suppression of the non-qualifying pixels across the image. The resulting image is cleaned from spurious noise, dilated and eroded and is finally tracked where upon our tracker returns pixel co-ordinates corresponding to the centerlines of the vessels. Once the centerlines are identified we move on to determine size measurements such as length, width of the vessels and their respective walls. We propose and implement as measure of tortuosity of a vessel the average angle of turn, sampled at appropriate steps. All measurements are calibrated for magnification errors with the help of the real world papilla size which we consider a known invariant. The papilla is located automatically with a disk fitting procedure that exploits very minimum assumptions. Throughout our work there is no dependence on any kind of threshold or parameter fine tuning.

6.3 Evaluation of the work

Our algorithm for the automatic location of the papilla works well and is quite robust in the sense that reasonable deviations from the true centre of the feature do not make the program misquoted the radius. It downpresents a problem of time, but this is an option into consideration when one has to pick a final implementation. Also, there is an apparent trade off between the time and the accuracy of the estimate. If one had to re-implement it, we feel that the method would benefit from a compiled computer language. The non-maximal suppression works fine as well, bringing enough information to the next stages. Unfortunately, this information is not utilised fully by the tracker which produces fragmented vessels. Although our tracker identifies a good 80% of what can be seen with the human eye (arteriovenule subdivisions up to the fourth level) we feel that it would benefit greatly from a post-processing linking routine. Linking should be as “intelligent” as possible and should keep the interpolation of points very restricted. Also, this

way we could possibly look under a different, more global perspective to length and tortuosity. The whole program would greatly benefit from extended testing as time did not permit rigorous testing at this point. It is well known how biological variation makes comparisons very difficult.

6.4 Future work

This project was meant to provide the software basis for an extended set of experimental measurements concerning a possible correlation of results that could lead to conclusions on the relationship between retinal findings and hypertension. In that sense, we feel that the ridge detection technique has proved ideal for our purposes and could well be used in an extension of this work into an integrated software tool in support of a physician's practise. A point which someone should definitely look into is the tracker, which should either be redesigned and tailored to our needs or complemented with a powerful linking engine. Further extensions would be to work with coloured images as well. Coloured images could be reduced to gray scale for the purposes of segmentation and then one could go back to the colour image without any registration problem. This way one could study standard arteriovenous indices such as artery to vein ratio, or arteriovenous crossings.

If the techniques presented here prove to be robust and reliable under strong experimentation as we feel confident they will, a further step would be to form a sort of grading scale on the severity of hypertension. Ideally, classification in this scale will be by the software itself either through the use of classic AI techniques, e.g. through an Expert System, or through a Connectionist approach where a Neural Network could be trained to recognise the hypertensive patient's symptoms.

Lastly, we would wish to point out that the use of computer technology as a diagnostic aide should not be seen isolated from the other uses of automation in use today by many people in the medical community. Such software should most probably be able to connect directly, or through a custom interface, to a patients' database. Such a move would indicate excellent use of resources as well as foretell the integrated role computers are going to play in medicine within the next few years.

Bibliography

- [Baumbach *et al.* 89] P. Baumbach, Rassow B., and Wesserman W. Absolute fundus dimensions measured by multiple beam interference fringes. *Invest Ophthalmology Vis Sci*, 30:2314–2330, 1989.
- [Beattie 82] R. Beattie. *Edge Detection for semantically based early vision*. Unpublished PhD thesis, University of Edinburgh, 1982.
- [Bennet *et al.* 94] AG Bennet, Rudnicka AR, and Edgar DR. Improvements on littman’s method of determining the size of retinal features by fundus photography. *Graefe’s Arch Clin Exp Ophthalmol*, 232:361–367, 1994.
- [Canny 86] John Canny. A computational approach to edge detection. *IEEE Trans. PAMI*, 1986.
- [Fishler & Wolf 83] MA Fishler and H. Wolf. Machine perception of linear structure. In *Proceedings IJCAI*, 1983.
- [Jagoe *et al.* 90] R. Jagoe, Smith P, Blauth C, Taylor KM, and Wooton R. Automatic geometrical registration of fluorescein retinal angiograms. *Computers and Biomedical Research*, 23:403–409, 1990.
- [Jagoe *et al.* 93] R Jagoe, Arnold J, Blauth C, Smith LC, Taylor KM, and Wooton R. Retinal vessel circulation

- patterns from a sequence of computer-aligned angiograms. *Investigative Ophthalmology and Visual Science*, 34(10):2881–2887, 1993.
- [Keith *et al.* 39] N.M Keith, H.P Wagener, and N.W Barker. Some different types of essential hypertension: Their cause and prognosis. *Amer. J. Med Sci*, 197(332), 1939.
- [Li *et al.* 89] Du Li, G.D. Sullivan, and K.D Baker. Edge detection at junctions. In *Proceedings AVC*. Fifth Alvey Vision Conference, 1989.
- [Littmann 82] H. Littmann. Zur bestimmung der wahren grosse eines objektes auf dem hintergrund des lebenden auges. *Klin Monatbl Augenheilkd*, 180:226–289, 1982.
- [McInerney & Terzopoulos 95] T McInerney and Demetri Terzopoulos. Topologically adaptable snakes. In *IEEE International Conference on Computer Vision*, MIT, Boston, 1995.
- [Miles & Nuttal 93] FP. Miles and AL. Nuttal. Matched filter estimation of serial blood vessel diameters from video images. *IEEE Transactions on Medical Imaging*, 12(2):147–152, 1993.
- [Sanchez A 95] Honrubia FM Sanchez A, Larrosa JM. Improvements on littman’s theoretical method. *Graefe’s Arch Clin Exp Ophthalmol*, 233(8), 1995.
- [Stokoe 77] N.L. Stokoe. Arteriosclerosis and hypertension. In A.T. Proudfoot, editor, *The Eye in Medicine*, volume 49, pages 39–55. The Royal College of Physicians of Edinburgh, 1977.

- [Szolovits(ed) 82] P. Szolovits(ed). *Artificial Intelligence in Medicine*. Westview Press, Boulder, Colorado, 1982.
- [Vink 70] R. Vink. *Fluorescein Angiography in Diabetic Retinopathy*. Leiden, 1970.
- [Williams 84] R. Williams. *Medical Photography Study Guide /4ed*. Lancaster, MTP Press, 1984.
- [Wise et al. 71] G. Wise, Dollery C, and Henkind P. *The Retinal Circulation*. Harper and Row, 1971.
- [Zhou et al. 94] L. Zhou, Rzeszotarski MS, Singerman LJ, and Chokreff JM. The detection and quantification of retinopathy using digital angiograms. *IEEE Transactions on Medical Imaging*, 13(4):619-626, 1994.

Appendix A

Experimental Results

In the following pages we present every step of the method described above applied on our test image for demonstration purposes. After smoothing and building the directional map (Figure 4-12) we move on to suppressing non-ridge pixels. The outcome of scanning the image through the three criteria presented earlier can be seen in Figure A-1.

Next comes cleaning individual pixels which can be seen in Figure A-2 as well as dilation (Figure A-3) and erosion (Figure A-4).

Figure A-5 depicts the results of the two thinning filters and figure A-6 the results of the tracking algorithm. Lastly, all tracks with length smaller or equal to 10 are removed. The final track file can be viewed in Figure A-7.

It should be noted that all track files have been turned to HIPS format using

```
tretob -h < file.trk > file.hips
```




Figure A-1: Non-maximal Suppression

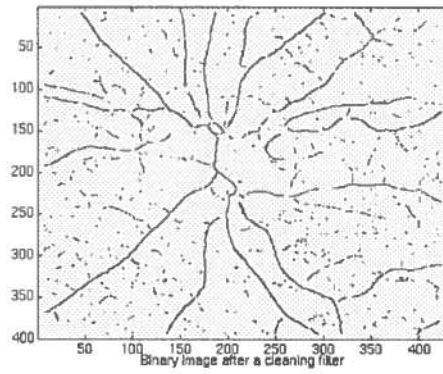


Figure A-2: Cleaned image




Figure A-3: Dilated image




Figure A-4: Eroded image




Figure A-5: thinned image




Figure A-6: Before removal of small tracks




Figure A-7: After removing small tracks