COMPUTER VISION ALGORITHMS FOR RETINAL VESSEL WIDTH CHANGE DETECTION AND QUANTIFICATION

A proposal to conduct doctoral study at Rensselaer Polytechnic Institute

Kenneth H. Fritzsche
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
Email: fritzk2@rpi.edu
Phone: 276-6909

Thesis Advisor
Charles V. Stewart
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
Email: stewart@rpi.edu
Phone: 276-6731

Thesis Advisor
Badrinath Roysam
Department of Electrical, Computer, and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
Email: roysab@rpi.edu
Phone: 276-8067

October 2002

Contents

1 Introduction

The proposed work is motivated by the needs of the medical community to accurately and consistently identify and quantify changes in blood vessels in the retina over time. This need arises from the numerous eye and systemic diseases that affect the vessels of the retina and can be diagnosed by detecting change in retina vessels.

1.1 Retina Vessel Change

There are multiple eye diseases that affect the vasculature in the eye, particularly in the vessels of the retina. These diseases can cause physical changes to existing vessels, such as changes in the width, color, and path of the vessels. They can also cause neovascularization — the growth of new vessels. Table 1 shows how some changes in vessels are linked to different eye diseases.

	Manifestation							
Disease	NVE	NVD	Artery Color	Vein Color	Artery Narrowing	Artery Dilation	Vein Narrowing	Vein Dilation
Choroidal Neovascularization	x							
Diabetic Retinopathy	x	x	x		x		x	x
Hypertensive Retinopathy				x	x			
Branch Retinal Vein Occlusion		x		x	x	x		x
Central Retinal Vein Occlusion		x		x	x	x		x
Hemi-Central Retinal Vein Occlusion				x	x			x
Branch Retinal Artery Occlusion		x	x		x		x	
Cilio-Retinal Artery Occlusion		x			x		x	
Arteriosclerotic Retinopathy			x					
Coats		x		x	x	x	x	x
Macro-aneurysm					x			
Hollenhorst Emboli					x			

Table 1: The above table shows different manifestations of disease that affect the blood vessels of the retina. (NVD = Neovascularization at the optic disk, NVE = Neovascularization elsewhere).

While many eye diseases alter the blood vessels in the retina, there are also systemic diseases that affect the blood vessels in the entire body. Since the eye is the only place for a physician to directly observe blood vessels in vivo, the retina vasculature offers a unique

opportunity for a physician to gain valuable clues for detecting and diagnosing systemic disease. For instance, hypertension increases large artery dilation in the body by as much as 15% [38] but in the eye it has been found to increase retinal artery dilation by as much as 35% [18]. Age and hypertension are thought to cause changes in the bifurcation geometry of retinal vessels [36]. Additionally, retinal arteriolar narrowing is thought to precede the onset of diabetes [43] and in women is related to the risk of coronary heart disease [42].

1.2 System Requirements

Prior work in the research literature consider global properties of the retina. This requires computing summary descriptions, such as the average width, of the retinal vasculature in a single image and comparing the resulting numbers against previous numbers for the same individual or comparing the numbers against population distributions. These yield coarse measures that require large-scale changes to be significant.

We propose a radically different approach: detecting longitudinal changes at the level of individual blood vessels or individual locations on individual blood vessels. This is enabled by the recently developed Dual-Bootstrap ICP algorithm [37] which is able to precisely align virtually any pair of images taken of the same retina, even if these images are taken months apart and show substantial, disease-induced changes in the retina. Such precise registration opens an entirely new realm of possibilities for investigating disease-induced changes. Our goal in this thesis is to develop the tools to exploit registration in the detection of vessel change (such as those seen in Figure 1) in order to direct physicians attention to possible problematic areas.

Detecting changes in registered images of individual vessels requires a number of capabilities in addition to registration. Algorithms are needed to extract the vasculature, measure its properties, establish correspondence between vessels in aligned images, measure change, and determine the significance of the changes. Several of these algorithms

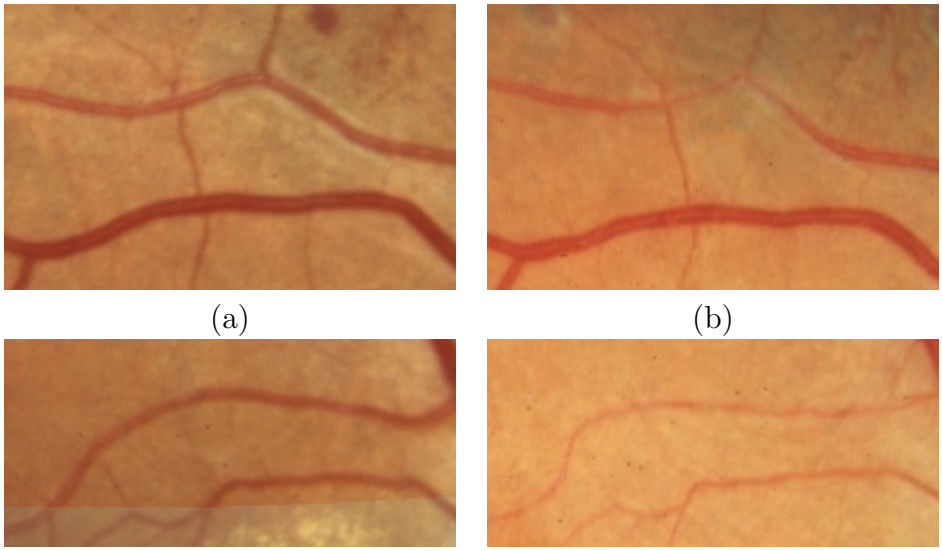

Figure 1: Illustrating changes in vessel caused by disease. All four images are from the same patient. Images (b) and (d) were taken 18 months after images (a) and (c). Note the pronounced thinning of the vessels in (b) and (d) as compared with the same vessels in (a) and (c). Also note the large shift in the vessels in image (d) at the point where the top and bottom vessel bends (in the left half of the image).

already exist, as illustrated in Figure 2, although all require improvement and must be validated for the change detection application. These form the outline of our current and proposed work. We break them up into the following categories:

Vascular modeling and feature extraction: A variety of vessel models have been proposed in the literature (see Section 2.3). I believe a simple parallel-edge model is the most useful. Already, I have made substantial contributions to improving a parallel-edge retina vessel tracing algorithm developed originally by Can [4]. I propose to further improve this algorithm using a twin-sided, ribbon snakes technique to obtain more accurate boundaries and then experimentally analyze the parallel-edge model by comparing to other models.

Validation Change-detection requires reliability in extracting vessels and measuring their properties. I propose to develop algorithms to measure the success of techniques

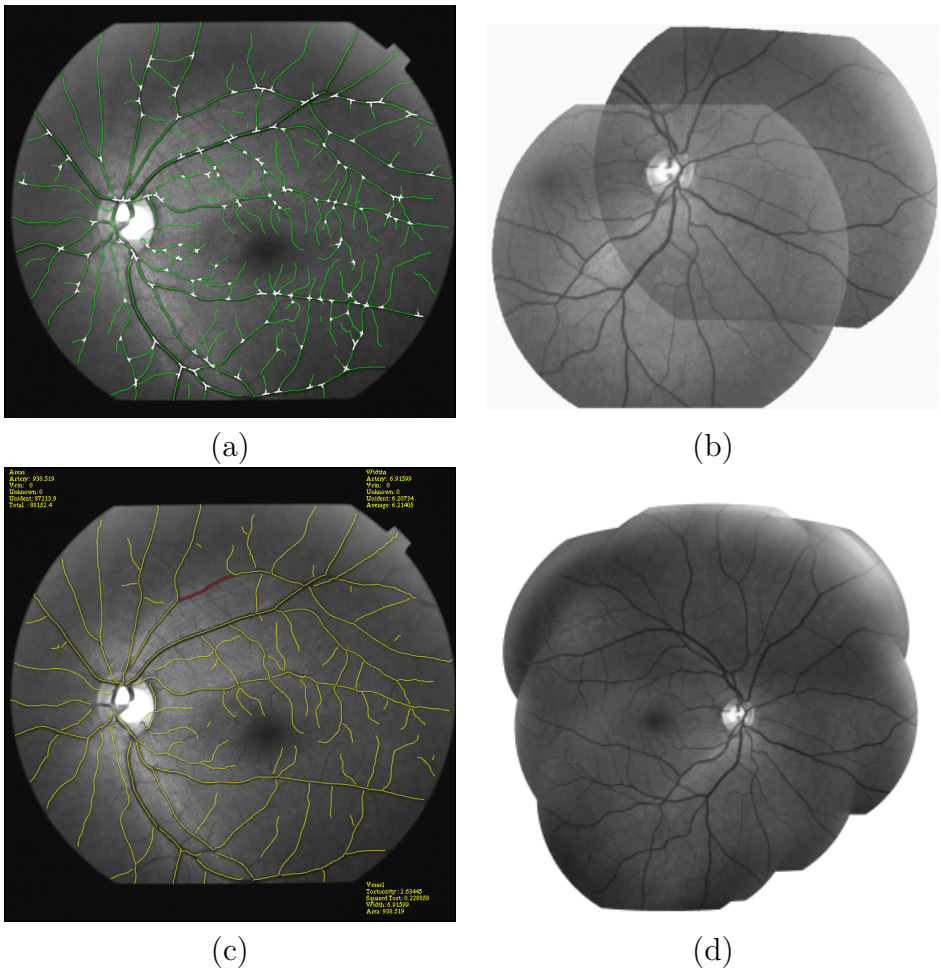

Figure 2: Image (a) illustrates the results of vessel extraction. Image (b) illustrates the alignment of two individual images. Image (c) illustrates the ability to calculate vessel morphometries. The text in the top corners of the image show aggregate vessel statistics and the bottom right shows information for the selected vessel. Image (d) demonstrates the ability of the registration algorithm to align multiple images and build a single "mosaic" image.

both as compared to (approximate) ground truth and in terms of repeatability of measurements.

Multi-image feature extraction. Frequently, multiple images are taken of the same retina during the same sitting, but from different viewpoints. Aligning these images allows for the collection of redundant measurements. This redundancy makes it possible to overcome the effects of varying illumination and glare. Thus, I propose to use the redundant measures to develop a multi-image vessel extraction method to generate a more complete extraction that is consistent between all images.

Change detection: All of the foregoing are precursors to actually detecting changes. The actual change detection algorithm will include methods for establishing correspondence between different images of the same vessel, distinguishing between vessels that are truly missing in the images and just missed by the algorithms, calculating change, and measuring its significance.

All of these algorithms will be combined into a C++ software system and user interface.

These research areas will be discussed in the remainder of this proposal. Sections 2 through 5 define each topic in more detail. Each section starts by presenting a single sentence mission statement of what is to be accomplished in the specific area. The section then goes on to discuss pertinent background material and/or achievements by other researchers in the area. The remaining part of the section provides a proposed methodology with which I intend to accomplish the stated research mission. Section 7 will list the expected significant contributions resulting from my work. Section 8 provides a proposed time line for completion of the work. Section 9 shows my current assessment of completion for each part of my research.

2 Single Image Vessel Extraction and Description

2.1 Mission

To extract and describe the blood vessels that appear in a single retinal image.

2.2 Discussion

Separating the portions of a retinal image that are vessels from the rest of the image (or background) is known as vessel extraction or vessel segmentation. These processes are similar in goal but distinct in process. Both are inherently hard due to such factors as poor contrast between vessels and background; presence of noise; varying levels of illumination and contrast across the image; and physical inconsistencies of vessels such as central reflex [14, 3, 33] and pathologies. Algorithms have been developed to address these issues but multiple images taken from the same patient in a single session still yield different segmentation/extraction results. This section will present some models used to identify blood vessels, discuss some of the different segmentation/extraction techniques, present work that I have completed thus far, and propose a methodology to improve extraction results.

2.3 Vessel Models

There are varying models used to identify vessels in images. Most are based on detectable image features such as edges, cross sectional profiles or regions of uniform intensity. Edge models work by identifying vessel boundaries typically by applying an edge detection operator such as differential or gradient operators [32, 40, 4]; Sobel operators [41, 35]; Kirsch operators [26]; or first order gaussian filters [25]. Cross sectional models try to find regions of the image that closely approximate a predetermined shape such as a half ellipse [22]; gaussian [46, 6, 17, 14, 8]; or 6th degree polynomial [16]. Algorithms that

use local areas of uniform intensity generally employ thresholding [20], relaxation [1], or morphological operators [21, 17, 44, 45].

For vessel change detection, the most important aspect of vessel segmentation is the ability to accurately and precisely locate vessel boundaries. Of these techniques, it appears that the most appropriate for accurate location of vessel boundaries are the methodologies that are boundary detection based. Cross sectional and intensity region models results typically produce images that require further processing to extract vessel boundary information. These produced images can be sensitive to selected thresholds and the boundaries extracted may be slightly modified from the original based on parameters used to construct them (such as the value for σ when using a gaussian profile model). Also defining the appropriate definition of a boundary for these methods is a source of concern (which does not exist for boundary detection models) and will be discussed further in Section 5.4.3. These methods need to be examined and I intend to do a repeatability study as described in Section 5.5 to ascertain that boundary models are most appropriate for vessel width change detection.

2.4 Previous Vessel Extraction Methods

Algorithms for identifying blood vessels in a retina image generally fall into two classes — those that segment vessel pixels and those that extract vessel information. Generally segmentation is referred to as a process in which, for all pixels, certain characteristics of each pixel and its neighbors are examined. Then, based on some criteria, each pixel is determined to belong to one of several groups or categories. In retinal image segmentation, the simplest set of categories is binary — vessel or non-vessel (background). This is what we refer to as "vessel segmentation". Techniques for segmentation include using morphological operations [21, 17, 44, 45], matched filters [6, 41, 17, 8, 12], neural nets [23, 34], FFTs [39] and edge detectors [41, 26, 32] to produce binary segmentations. Seg-

mentation algorithms are generally computationally expensive and require further analysis to calculate morphometric information such as vessel width.

Extraction algorithms [25, 14, 31, 4, 46] on the other hand, generally use exploratory techniques. They are faster computationally and usually determine useful morphometric information as part of the discovery process. They work by starting on known vessel points and "tracing" the vasculature structure in the image based on probing for certain vessel characteristics such as the vessel boundaries. Since the vessel boundaries are part of the discovery process, these algorithms generally contain information such as vessel widths, center points and local orientation that can be used later in the change detection process. For these reasons, I decided to an exploratory algorithm in my work and selected Can's algorithm [4] as the starting point of my research.

2.4.1 Can's Vessel Extraction Algorithm

Can's vessel extraction (or "tracing") algorithm [4] uses an iterative process of tracing the vasculature based on a localized model. The model is based on two physical properties of vessels - that vessels are locally straight and consist of two parallel sides. As such, it works by searching for two parallel vessel boundaries, found based on detection of two parallel, locally straight edges. The entire algorithm consists of two stages.

Stage 1 (seed point initialization): The algorithm analyzes the image along a coarse grid to gather gray-scale statistics (contrast and brightness levels) and to detect initial "seed" locations on blood vessels using the gray-scale minima. False seed points are filtered out by testing for the existence of a pair of sufficiently strong parallel edges around the minima. The initial seed point is filtered out if the two strongest edges do not both exceed a contrast threshold (generated from the gray-scale statistics) or if the directions of the two strongest edges are not sufficiently similar (within 22.5 degrees). On average, about 40% of the initial points are filtered out by this procedure.

Stage 2 (recursive tracing): The second stage is a sequence of recursive tracing steps

that are initiated at each of the filtered seed points. For each filtered seed, the tracing process proceeds along vessel centerlines by searching for vessel edges. These edges are searched from a known vessel center point in a direction orthogonal to the local vessel direction to a distance of one half the maximum expected vessel width. (Note all directions are from a discrete set of angles, usually 16 angles, with an interval of 22.5 degrees). Three angles are searched (the previous center point direction $\pm$ 1) and the angle of the strongest edge is estimated. The next trace point is determined by taking a step from the current trace point in the direction of the strongest edge. The resulting point is then refined by applying a correction that is calculated based upon how far the new point is from the center of the points at which the left and right boundaries are found. This new point is kept only if it does not intersect any previously detected vessels, is not outside the image frame, and if the sum of edges strengths is greater than the global contrast threshold calculated during seed point detection.

Can's algorithm's main purpose was to extract vessels to identify landmarks (vessel bifurcation and cross-over points) as a basis for registration. While effective at producing landmarks, vessel boundaries were often not smooth and vessel segments were often missed.

2.5 Work Done So Far

In order to overcome some of the limitations in Can's algorithm, I implemented a modified version of Can's algorithm using an open source, multi-platform, computer vision and understanding collection of C++ libraries called VXL. The modified algorithm is implemented in both a command line and GUI version with results being generated in both text and in any of multiple image formats.

2.5.1 Smoothing Vessel Boundaries

As discussed in Section 2.4.1, in order to find the vessel boundaries, Can applied templates iteratively from a point known to be a vessel center point outwards to some distance $M/2$ where M is the expected maximum width of a blood vessel. This can be seen in Figure 3. This technique allows vessel boundaries for adjacent vessel segments to vary significantly. In an attempt to produce smoother boundaries, we modified this search strategy to only search for a vessel edge at a predicted edge location $\pm d/2$ where d is the maximum allowable change in width of a vessel at each step of the tracing algorithm. Taking into account that there are two boundaries, at each step of the algorithm, the largest possible difference in width between adjacent vessel segments is d. Note that the search space is also constrained in that the predicted edge location plus $d/2$ must be less than $M/2$ and the predicted edge location minus $d/2$ must be on the correct side of the predicted center.

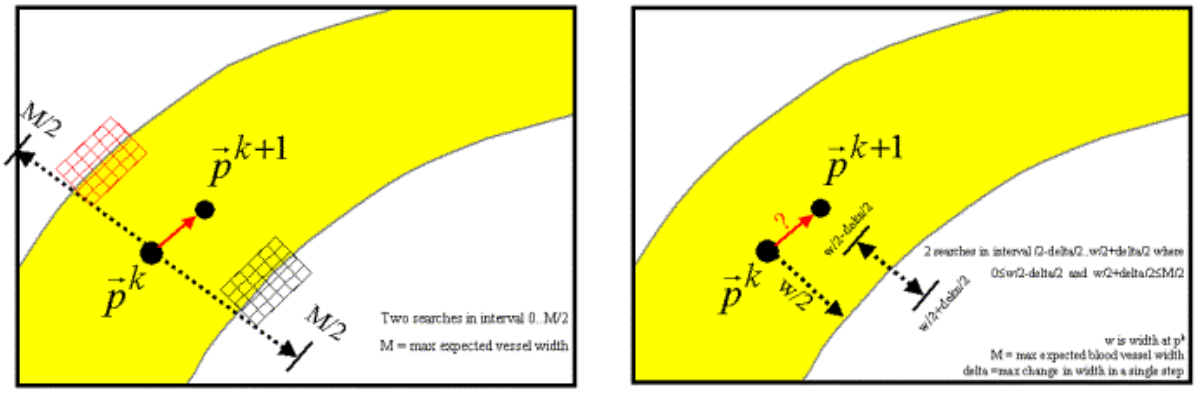

Figure 3: Illustrating Can's original search strategy for vessel boundaries (on the left) and the modified search strategy to provide smoother boundaries (on the right). While computationally more efficient and successful at providing smoother boundaries, the modified method still allows too much variance over the length of the entire vessel.

2.5.2 Other Modifications to Can's Algorithm

The next modification of Can's base algorithm is the implementation of local thresholds. Can's original algorithm uses a single global threshold to determine vessel boundaries. This single global threshold resulted in missed vessels, particularly in images with a large degree of variation caused by uneven illumination or the presence of pathologies. By breaking the image into smaller regions and calculating thresholds for each region, more appropriate thresholds were calculated and better results achieved.

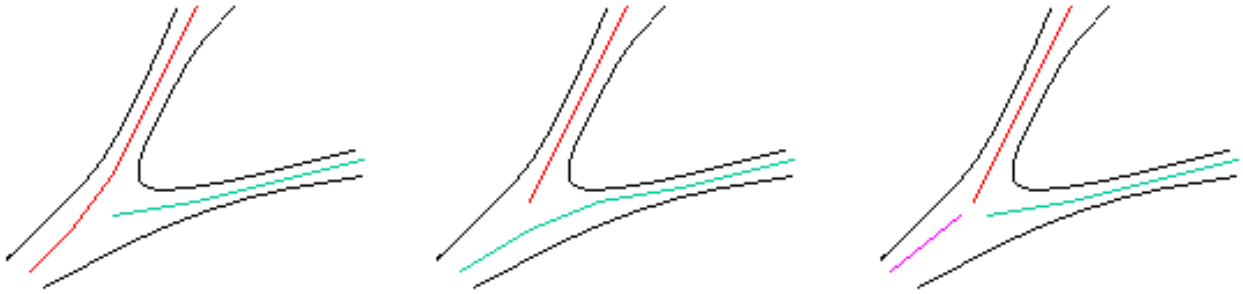

Figure 4: The left and middle image demonstrate possible results of the original algorithm in two different images. It would be difficult to compare the three recognizable vessel segments as each set of results contains only two segments that differ greatly. The image on the right depicts how the current algorithm, would always represent the results as three separate segments in all images and as such would be much easier to compare.

A third modification concerns breaking vessels into segments. In order to compare corresponding vessel segments between images, it is necessary to present the results of the vessel extraction algorithm in a manner that facilitates identification of corresponding segments. To accomplish this, a simple graph structure of the vasculature is built by treating all landmarks (i.e. bifurcation and cross over points) as nodes and all traces as edges. Vessel segments are defined as starting and ending at either a landmark or at nothing. This requires that detected vessels be split into multiple segments (i.e. graph edges) at each landmark. While this seems to be a trivial distinction, the original algorithm was not concerned with building a vessel graph or with identifying segments. As such it

made no effort to break a trace into individual segments. So, in the original algorithm, it was possible for a vessel that contained a single bifurcation to be represented by only two traces. In the modified version, it is now always represented by three. Figure 4 illustrates this concept.

Still other modifications improving the accuracy or efficiency of the tracing algorithm were implemented. These are not vital to the proposed work and warrant no discussion other that to mention them as follows:

1. better initial seed point determination;
2. use of multiple lengths of templates;
3. use of a binning structure to cache results;
4. interpolation of results to sub-pixel accuracy;
5. better accuracy and repeatability in determining landmark location;
6. inclusion of a momentum strategy to allow tracing to trace through noisy parts of an image.

2.5.3 Limitations of the Modified Can Algorithms

While the above improvements have made an improvement to Can's original algorithm, there are still some limitations, particularly in two areas. First is in the area of detecting extremely small vessels, particularly neovascular vessels. This is primarily caused by lower contrast of the thin vessels (typically 1-2 pixels in width) and their highly tortuous nature. Thus the locally straight parallel edge model fails and a new method for finding neovascular vessels is needed. However, this is not a part of my proposed work.

The second limitation is in the area of precise location of the vessel boundaries. This is vital for detection of changes in vessel width. It is influenced by factors in the image (such

as poor contrast and noise) and by the methodology used to estimate vessel boundaries. Inaccuracies are caused by modelling the boundary with a set of boundary points found by using a discrete set of edge detectors oriented over a discrete set of directions. Additionally, poor contrast and noise result in yet more error. These types of errors are illustrated in the results shown in Figure 5. This degree of error may be acceptable in certain applications such as low-resolution segmentation or feature extraction for registration, but it is not tolerable at the level of resolution desired for change detection. A better method for determining vessel boundaries is is needed.

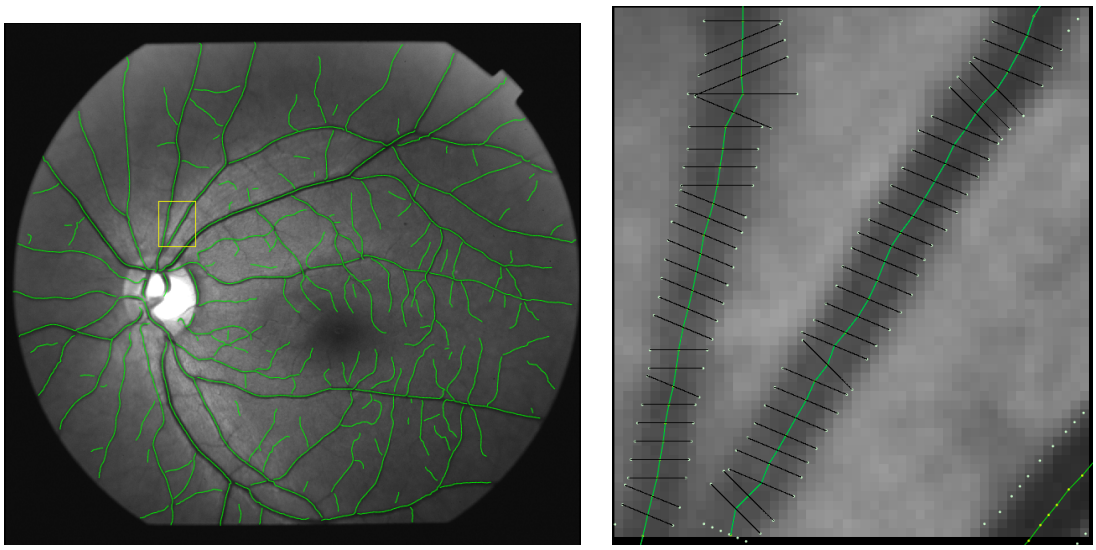

Figure 5: Illustrating the inaccuracy in the boundaries. The figure on the left shows an image with a fairly accurate vessel (centerline) segmentation. The image on the right, shows in detail the boxed region from the image on the left. Note the inaccuracies in the vessel boundary points caused by discretization errors or noise. Also note how for each trace point, the corresponding boundaries are often not perpendicular to the orientation of the vessel and as such are not accurate for use in measuring vessel width.

2.6 Proposed Methodology

Determining accurate vessel boundaries is essential for accurate measurement of vessel width and for a system to detect change in width over time. Both manual [9] and automated methods [12, 13, 31, 5] have been developed to measure vessel width. All are

based on first specifying and then measuring vessel cross-sections. However, these measurements can be prone to error depending on the method used to determine the start and end point of the cross section to be measured. In the case of the current modified version of Can's algorithm, the boundaries and cross sections are defined by a set of boundary points. These boundary points represent the center of a detected locally straight (i.e. typically 5, 9 or 13 pixels in length) edge. These edges are found from a search that uses a discrete set of angular orientations and are often inaccurate. Additionally, the resulting cross section formed from corresponding boundary points are not truly perpendicular to the local orientation of the blood vessel as demonstrated in Figure 5. To eliminate these inaccuracies, the boundaries need to be defined in an alternate fashion.

In defining a different vessel boundary model, I propose to start with the set of vessel boundary points as provided from the existing modified tracing algorithm. I then wish to utilize the original intensity information to find a more accurate fit for the vessel boundary. For this I propose to perform contour fitting based on snakes [24, 29], subject to the constraint that the opposite boundary is near-parallel. I plan to explore the use of ribbon snake models [27] where two parallel boundaries are discovered simultaneously for each segment on the snake. This will provide a smoother, more accurate boundary with which to measure and compare vessel widths. A brief discussion of Snakes and Ribbon Snakes follows.

2.6.1 Snakes

Snakes as originally introduced by Kass et al [24] are represented parametrically, with each position denoted as $v(s) = (x(s), y(s))$ and the snake being controlled by an energy function given by

$$E_{snake} = \int_0^1 E_{int}(v(s)) + E_{image}(v(s)) + E_{con}(v(s))ds, \tag{1}$$

where E_{int} denotes the internal energy of the snake; E_{image} represents external image forces for such characteristics in an image as lines, edges, and terminations of line segments and corners; and E_{con} is an external constraint force. E_{int} is commonly represented as

$$E_{int} = \frac{\alpha(s)|v_s(s)|^2 + \beta(s)|v_{ss}(s)|^2}{2} \tag{2}$$

where the first term controls the elasticity of the snake and the second term controls the rigidity. In the context of vessel boundary detection, E_{image} should be attracted to edges and is represented as $E_{image} = -|\nabla I(x,y)|^2$ and $E_{con} = 0$.

Using the set of detected boundary points for a single vessel boundary as the initial points, a snake can be used to solve for the final set of boundary points by minimizing the energy equation in Equation 2.6.1. This can be done by expressing and solving Equation 2.6.1 as the following set of Euler-Lagrange differential equations:

$$\frac{\partial x}{\partial t} = \frac{\partial E_{image}}{\partial x} - \frac{\partial}{\partial s}(\alpha(s)\frac{\partial x}{\partial s}) + \frac{\partial^2}{\partial s^2}(\beta(s)\frac{\partial^2 x}{\partial s^2}) \tag{3}$$

$$\frac{\partial y}{\partial t} = \frac{\partial E_{image}}{\partial y} - \frac{\partial}{\partial s}(\alpha(s)\frac{\partial y}{\partial s}) + \frac{\partial^2}{\partial s^2}(\beta(s)\frac{\partial^2 y}{\partial s^2}) \tag{4}$$

Converting to discrete space, these equations are solved iteratively until gradient descent convergence is achieved as approximated when the total change in v is below some threshold.

This method will yield a curve that is attracted to the boundary (as defined by the gradient). It is a better representation of the vessel boundary than the current vessel tracing code boundary points because the curve generated is smooth and is known to have converged on the best local boundary. This methodology can be applied to find both boundaries of a single vessel separately. However a more appropriate application of snakes to vessel extraction and accurate boundary location is described in the next section.

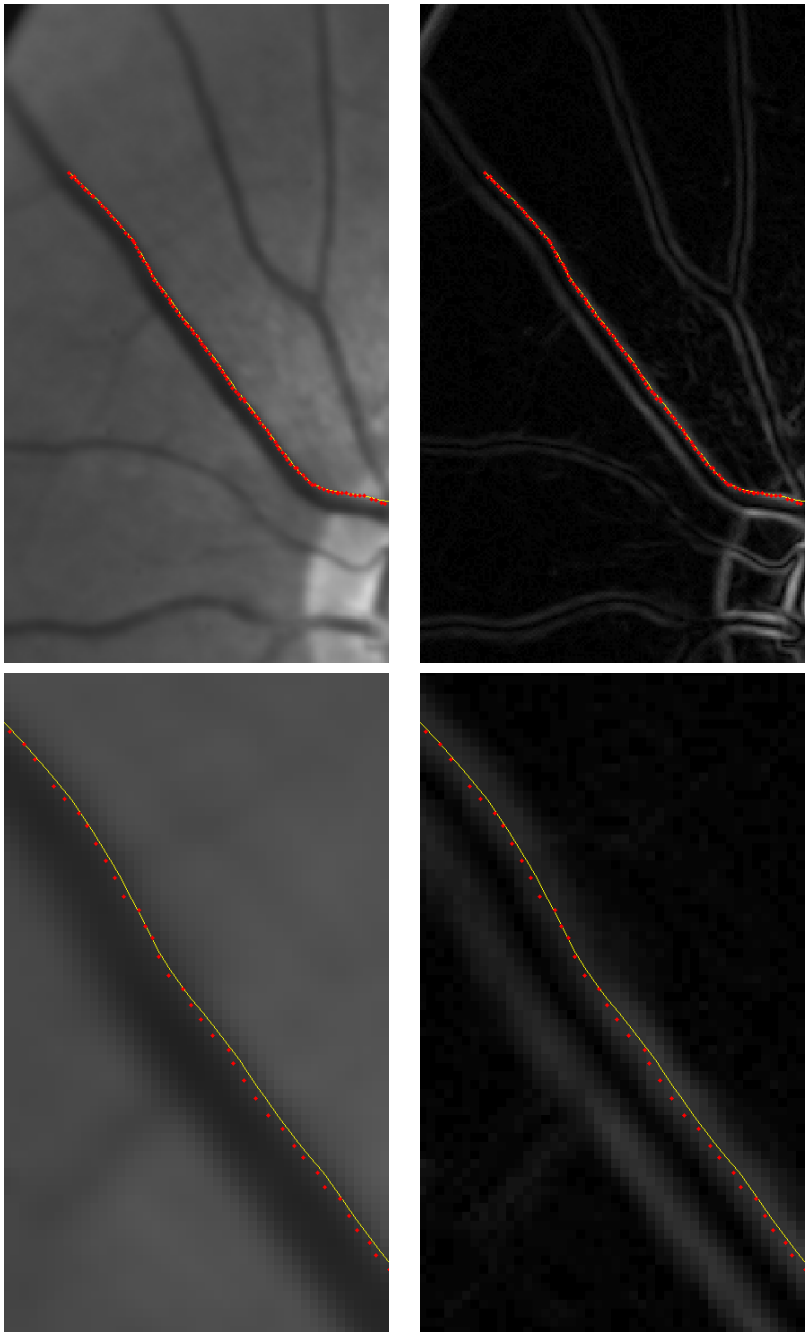

Figure 6: Illustrating the difference in boundary detection from vessel tracing algorithm and snakes. The top row shows the results of finding a vessel boundary using vessel tracing (red dots) and snakes(yellow line) on the original intensity image (left) and an image of the gradient magnitude (right). The bottom row shows a more detailed area of each. Note how the snake results are much smoother and differ by as much as 2 pixels when compared to the tracing algorithm's boundary.

2.6.2 Ribbon Snakes

Ribbon snakes can be used to detect linear features by locating the left and right boundaries and have been used to extract roads[27] in arial images. They are a better match for vessel boundary detection because they better model the properties of a vessel because they extract linear features that are bounded by two edges of contrast.

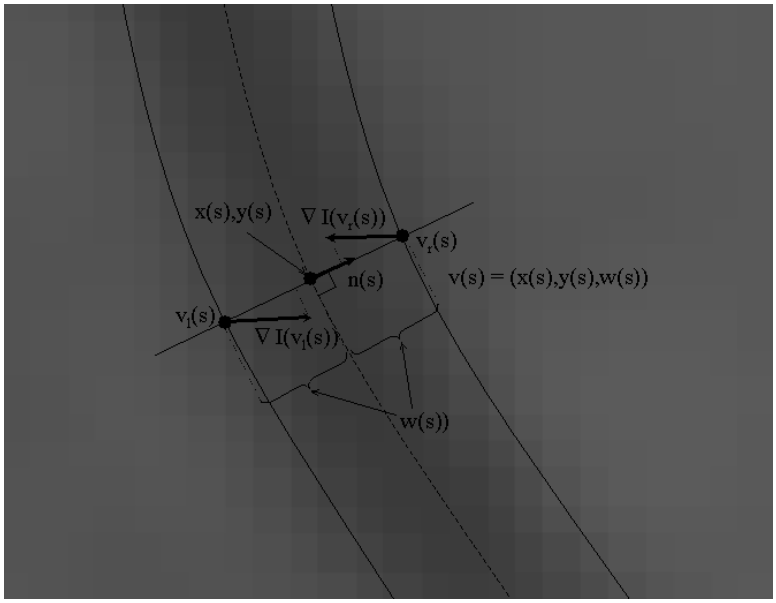

Figure 7: Illustration showing the parameters of a ribbon snake. Note the projection of the gradient on the unit norm ($n(s)$) used to further improve the boundary extraction.

Ribbon snakes extend the basic concept of a snake by introducing a third parametric component, namely a width component [11]. Thus each position is now represented as $v(s) = (x(s), y(s), w(s))$ where $x(s), y(s)$ represents the blood vessel center and $w(s)$ determines the vessel boundaries. Since we are trying to find two boundaries, E_{image} becomes

$$E_{image}(v(s)) = |(\nabla I(v_l(s))| + |\nabla I(v_r(s)))| \tag{5}$$

This energy term can be further modified to account for the knowledge that the vessel will be darker than the surrounding area. Thus on the left side of the ribbon snake, the

intensity goes from light background to dark vessel and on the right side of the ribbon, from dark to light. This can be represented by requiring the projection of the gradient vector onto the ribbon's unit normal vector to be positive on the ribbons left side and negative on its right as seen in Figure 7. This allow us to modify E_{image} as follows:

$$E_{image}(v(s)) = |(\nabla I(v_r(s))| - |\nabla I(v_l(s)))| \cdot n(s) \tag{6}$$

where $n(s)$ is the unit normal at $x(s), y(s)$ and $v_r(s)$ and $v_l(s)$ are the vessel right and left boundaries which can be expressed as

$$v_r(s) = w(s)n(s), v_l(s) = -w(s)n(s). \tag{7}$$

By supplying the centerline positions and widths for a vessel obtained from the modified Can's algorithm, the ribbon snake can be initialized and solved in much the same manner as normal snakes. The major difference is that the results will provide two contours representing both boundaries of the vessel. Additionally, ribbon snakes will allow me to explore the use of additional constraints to maintain the "almost parallel" boundaries by not allowing the width of the ribbon snake to vary greatly with respect to s.

2.6.3 Summary

Our enhanced implementation of Can's tracing algorithm still produces boundaries that are not sufficiently accurate for change detection. To overcome these inaccuracies, I will use snakes to perform a more robust boundary detection. The snakes will be initialized with the results from the modified tracing algorithm but the original image intensity structure will be again consulted and used to determine the final boundary locations. These boundaries should be more repeatable and more applicable for change detection.

3 Performance Measurement and Validation of Vessel Segmentation Algorithms

3.1 Mission

To validate results and measure the performance of vessel detection algorithms.

3.2 Discussion

The best method of validation is comparison against ground truth. Ground truth is a predetermined result that is known to be the "correct answer," and defines what *exactly* a computer algorithm is expected to produce. Unfortunately, this is often difficult or impossible to obtain [30]. This is the case for retinal images. Hence, methods must be developed to approximate the ground truth.

We focus much of the discussion on validation for segmentation algorithms because that is where the emphasis has been placed in the literature. Our own feature extraction algorithm can be compared to retinal image segmentation algorithms using this method. We then extend this methodology to develop validation for vascular centerline extraction. Finally, we develop a method for determining the repeatability of blood vessel width estimation.

3.3 Previous Methods

The arrival at ground truth is a known hard problem in image analysis and pattern recognition systems [19]. Previous research in fundus image vessel segmentation [22, 34, 23, 17, 32, 39] has often approximated ground truth by the creation of a human-generated segmentation also called a "gold standard" to which computer-generated segmentations are measured.

The generally accepted definition of a gold standard is a binary segmentation where each pixel assumes a value of 1 or 0 for vessel and background respectively. In comparing individual pixels of computer-generated results against the gold standard, there are four possible cases, true positive, true negative, false positive and false negative. These can be defined as follows. Let the gold standard be denoted as G and each individual pixel be denoted as $G(x, y)$. Likewise denote a computer generated segmentation as C and $C(x, y)$. Then,

$$C(x, y) = G(x, y) = 0 \longrightarrow \text{true negative;}$$
$$C(x, y) = G(x, y) = 1 \longrightarrow \text{true positive;}$$
$$C(x, y) = 1 \wedge G(x, y) = 0 \longrightarrow \text{false positive;}$$
$$C(x, y) = 0 \wedge G(x, y) = 1 \longrightarrow \text{false negative.}$$

The frequency of these cases provides data that can be used as an indication of an algorithm's performance.

Generation of gold standard images is often a costly and time-consuming process. In addition, it is known that human observers are subjective and prone to inter-observer and intra-observer variability. Thus the use of a single human expert's annotations is unreliable and should be considered inadequate for the purpose of generating a gold standard. Thus a more robust approach to the creation of a gold standard is to combine multiple human-generated manual segmentations. So, given a set of manual segmentations generated by multiple observers, H, with each individual segmentation being denoted H_i, we wish to obtain a single binary segmentation, G, that will be considered the gold standard. We know that all H_i segmentations will differ and thus a strategy for resolution of these differences must be created.

3.3.1 Creating Ground-truth from Conflicting Observers

In order to determine the best way to combine different observers segmentation results, one must first consider what constitutes a correct and incorrect vessel segmentation. Essentially, there are three possible ways to define correct vessel segmentation for a particular pixel. A conservative method is to declare pixels to be part of a vessel when all observers' segmentations agree. If such total agreement exists, then that pixel would be marked as vessel in the gold standard image. All other pixels would be considered background. In essence this amounts to computing a Boolean "AND" for all H_i in H. A less conservative method would be based on using majority rule. Each segmentation result would contribute a single vote for each pixel when determining if a pixel should be considered a vessel in the gold standard. If 50% or more, of the observers have determined a particular pixel to be a vessel, it would be marked a vessel in the gold standard. Note that the "majority" threshold could be set higher or lower as appropriate for the specific application. The least conservative method is to label a pixel in the gold standard as a vessel pixel if it is marked as a vessel in at least one of the observers' results. This last case is equivalent to computing the Boolean "OR" of the multiple observer segmentations.

3.3.2 Limitations of Previous Methods

Knowing that there is some uncertainty between observers, the above methods fail to capture the uncertainty in the gold standard. For instance, consider the case when only 1 in 10 observers determine a particular pixel to be a part of a blood vessel. There is really no way of knowing if this person is the only one who is correct or the only one who is wrong. Regardless of the approach used to develop the binary gold standard, the fact that only 10% of the observers thought that the pixel was a vessel is lost. So rather than deciding when the standard is being generated how many people it takes to be deemed "correct", this information should be captured in a modified gold standard in which each

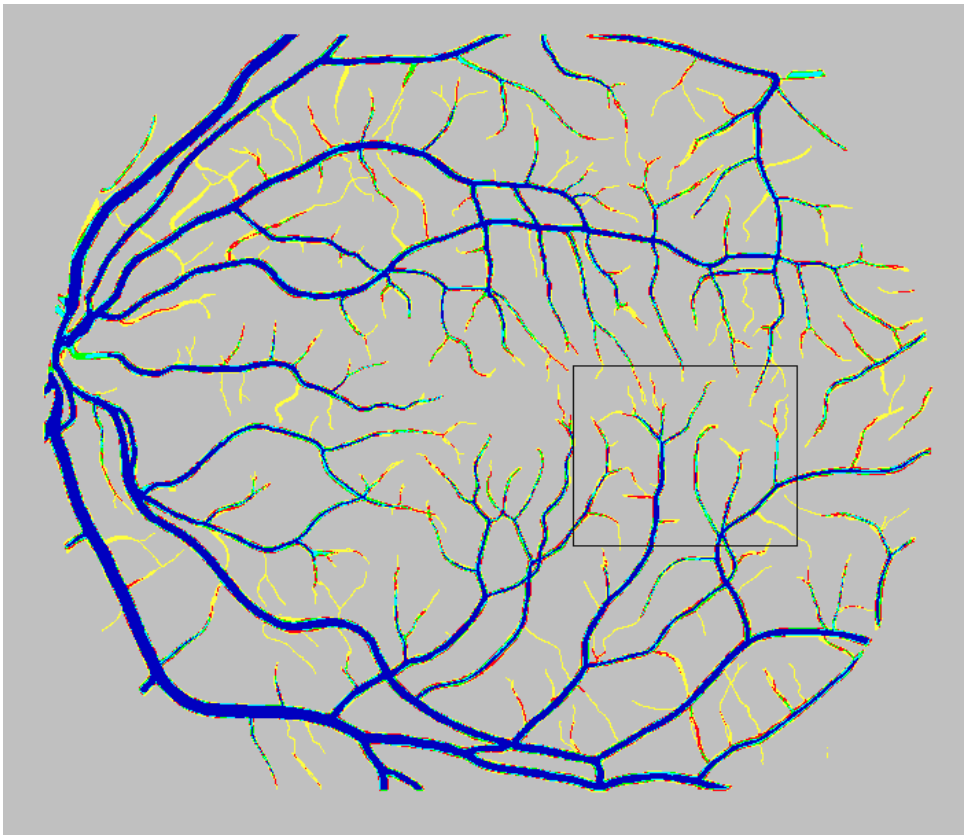

Figure 8: Multi-observer standard image formed from five separate hand tracings. Blue pixels are cases where five observers agree that the pixel is vessel. Cyan is four; green is three; red is two; yellow is one; grey is zero.

pixel is assigned a value based upon the number of observers who segmented that pixel as a vessel. This value can be considered a probability as described in the next section. Figure 8 illustrates this idea of using probabilities to build a non-binary gold standard formed from five separate hand traced images.

3.4 A Proposed Methodology

The driving ideas behind this methodology are threefold. First is to provide a standard that is representative of and maintains the uncertainty between the observers. Second is that when scoring an algorithm, mislabelled pixels should have a relative cost scale. For instance a mislabelled pixel that has total agreement of observers "costs" more than

those where there was not total agreement. Last is the idea that a computer segmentation not be penalized when its results agree with any observer. If the goal of a computer algorithm is to generate results in a manner comparable to a human, (though much faster and more consistent), an algorithm that demonstrates performance equal to that of any person should not be penalized for such behavior.

3.4.1 Validation Using a Probabilistic Gold Standard

Suppose that K is the number of observers ($K = 5$ in Figure 8). The segmentation label at each pixel for a computer-generated segmentation $C(x, y)$ can be thought of a binary random variable V from the set $\{0, 1\}$. A value of $V = 0$ is assigned at pixels labeled as background by an algorithm and $V = 1$ is assigned for pixels labeled foreground. Let M denote an integer valued random variable assuming values from the set $\{0, 1, ...K\}$ indicating the number of observers who have labelled a pixel as a vessel. The joint probability distribution $P_{VM}(v, m)$ can be estimated from the computer-generated segmentation and the manually scored segmentations. This distribution forms the basis for a performance measure S_{vessel} described below.

S_{vessel} is designed so that it assumes values in the interval from 0 to 1 where a 1 indicates a perfect score and zero a complete failure. A perfect score would indicate that all pixels determined by R or more observers to be vessel pixels have also been determined to be vessel pixels in C. A complete failure occurs when all vessel pixels determined by R or more observers are not labelled as vessel pixels in C. It is somewhere between 0 and 1 based on the value in the probabilistic gold standard for the pixels the computer does not label as vessel that are labelled as vessel by R or more observers. Note that in this method, described by the equation below, a missed vessel pixel in C, for which *all* observers agree to be a vessel pixel carries more weight in determining the final score than a pixel for which only R observers agree. R can be set to any appropriate value and we illustrate a "majority rules" idea in figure Figure 8 by choosing R such that it represents

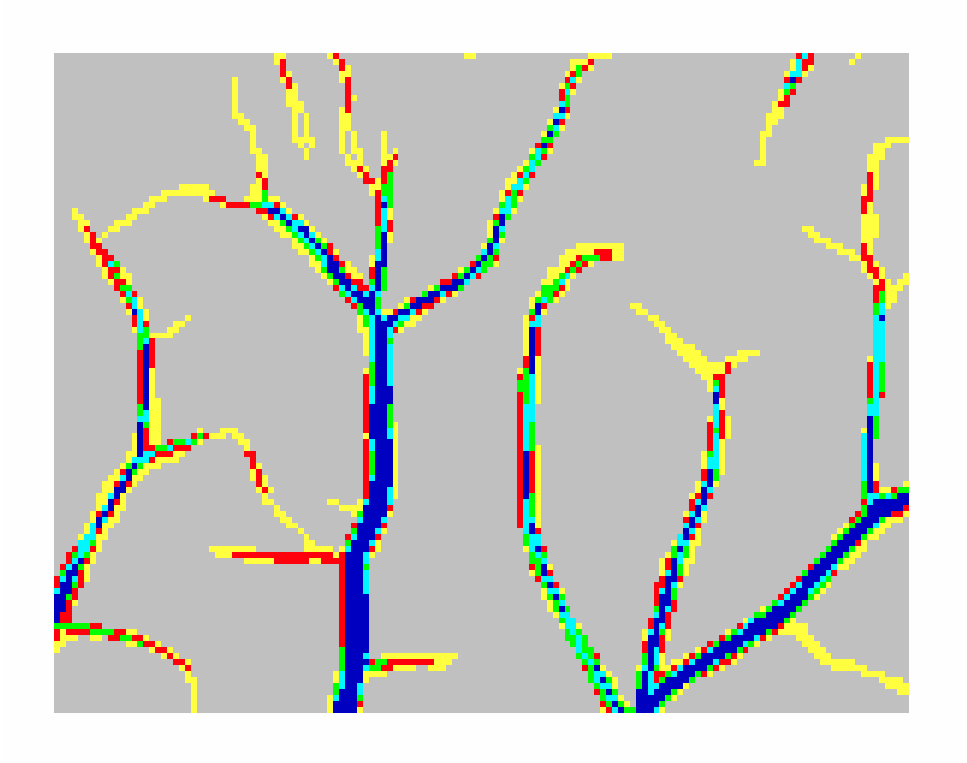

Figure 9: Enlarged portion of the boxed region in Figure 8 with the same color scheme described in Figure 8 caption.

a majority (50% or higher) of the observers ($R = \lceil K/2 \rceil$). Based on these considerations, we define the S_{vessel} as follows:

$$S_{vessel} = \frac{E[m \mid V = 1, M \geq R]}{E[M \mid M \geq R} \tag{8}$$

This is equivalent to

$$S_{vessel} = \frac{\sum_{m \geq R}^{K} mP[V = 1, M = m]}{\sum_{m \geq R}^{K} mP[M = m]}. \tag{9}$$

In the above equation, the numerator denotes the expected value of M for all the pixels for which the computer algorithm and R or more observers label as a vessel and can be thought of as the amount of agreement between C and the majority of observers. The

denominator represents the average value of M for all the pixels that R or more observers have labelled as a vessel and can be thought of as the "perfect score".

However, S_{vessel} does not take into account another indicator of segmentation performance, namely, the number of false positives which in our context is the number of pixels falsely identified as vessels. The following second measure account for this aspect of segmentation performance:

$$F_{vessel} = \frac{P\left[V = 1, M = 0\right]}{P\left[V = 1, M \geq R\right]}. \tag{10}$$

The numerator is the joint probability that the computer algorithm labelled the pixel to be a vessel and no observers label it as a vessel. The denominator is the probability that the computer algorithm labels it as a vessel, and more than R observers agree. In this measure, a score of zero means that there were no false positives. A score of 1 would indicate that an equal number of false positives to true positives was required to achieve the S_{vessel} score.

Using the same type of reasoning, it is possible to define similar metrics, denoted $S_{background}$ and $F_{background}$, respectively, for classification of non-vessel pixels as shown below:

$$S_{background} = \frac{\sum_{m<R}^{K} mP[V = 0, M = m]}{\sum_{m<R}^{K} mP[M = m]}. \tag{11}$$

and,

$$F_{background} = \frac{P(V = 0, M = K)}{P(V = 0, M < R)}. \tag{12}$$

Finally, the above metrics, $S_{vessels}$ and $S_{background}$ could be combined simply by averaging to develop an overall score between 0 and 1. The closer a set of the results is to the multi-observer standard the closer this score will be to 1. Figure 10 illustrates these

performance measures.

3.4.2 Validation for Results Generated by Vessel Extraction Algorithms

While the above measures would work for algorithms for which the goal is the segmentation of the entire vessel from the background, they would not work for algorithms that determine the centerlines of the vessels. For these algorithms, the following approaches can be used. The first approach is to convert a trace result into a segmentation result. This is often possible as most tracing algorithms determine the necessary data such as vessel width or boundary location for each point found on the trace. Thus a simple approach would be to segment as vessel points all the corresponding pixels between neighboring trace points based on the widths or vessel boundaries. This segmentation can then be scored as described above.

An alternative approach would be to test each point in the centerline trace against the multi-observer standard. Typically, exploratory algorithms generate results containing a set of vessel center points separated by a fixed or varying step size. So a simple procedure would be to test for true positives by seeing if a point in the centerline trace matches any non-zero (i.e. vessel) pixel in the multi-observer standard. If it does, it could be scored toward the S_{vessel} measure as described above or simply counted. Likewise false positives could be scored or counted and these combined measures can be used to compute measures indicating the algorithm's performance. However, such an approach fails to consider true and false negatives because it is debatable as to what exactly constitutes a true negative or false negative. Thus such an approach should not be considered and an approach that takes into account true and false negatives is discussed next.

A final approach is to build a trace based upon the multi-observer standard. This trace could be generated by running a tracing algorithm on the multi-observer standard, and if required, manually editing the results. These results, denoted G_{Trace} should then be used as the standard with which to compare tracing results. False/true positive and false/true

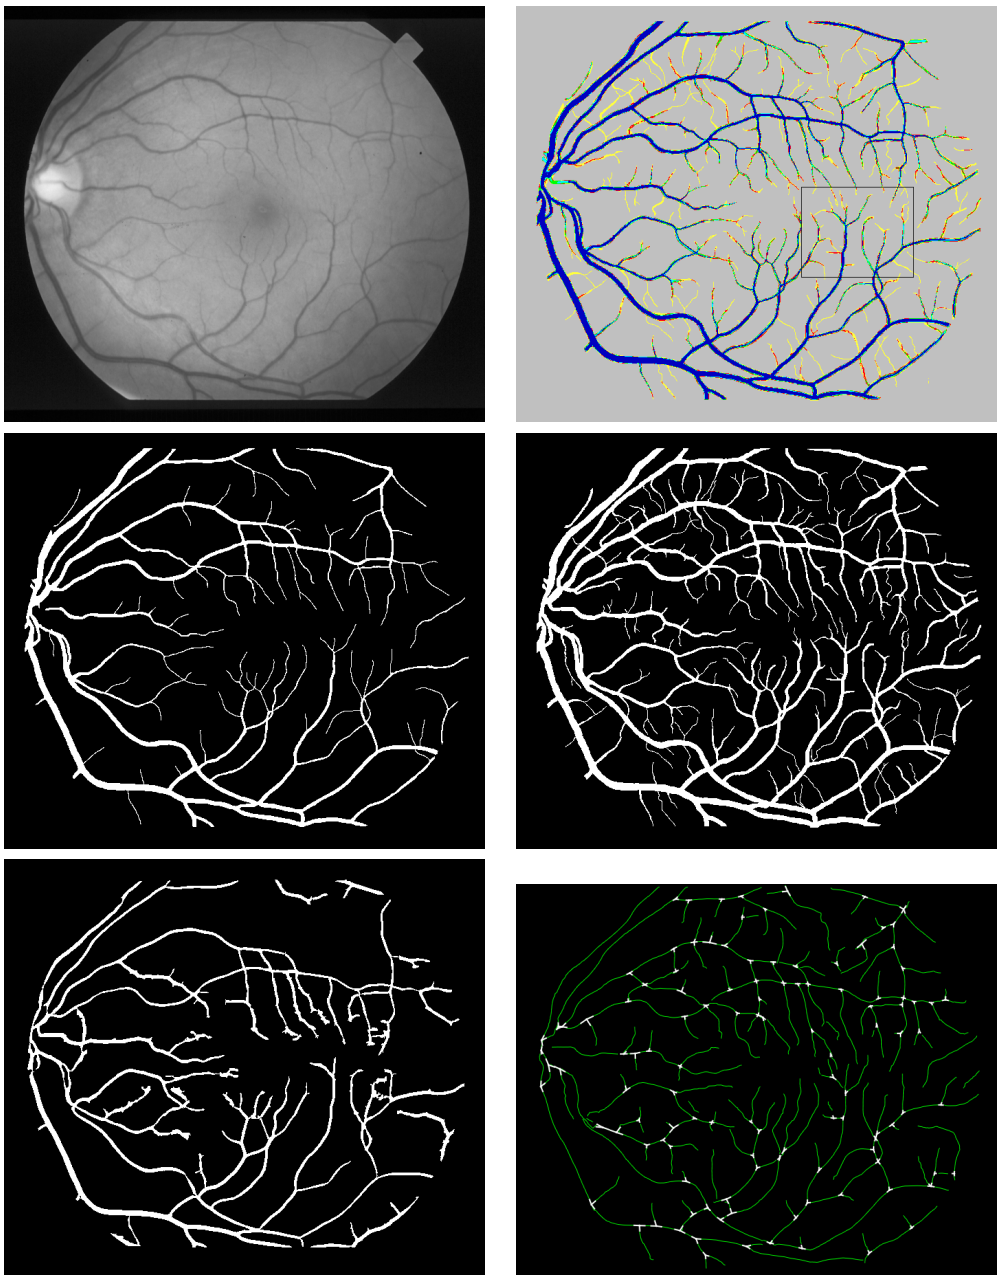

Figure 10: Illustrating the measurement of segmentation performance from multiple manual observers. Panel (a) is the original image. Panel (b) is the multi-observer gold standard generated from five hand traced images. Two of these manual traces are shown in Panels (c) and (d). Note the significant inter-observer variation. Panel (e) is a segmentation produced by Hoover's algorithm [17]. For this result, $S_{vessel} = 0.69$, $F_{vessel} = 0.14$, $S_{background} = 0.98$, and $F_{background} = 0.01$. Panel (f) is a tracing generated by the modified Can's algorithm [4]. Before scoring, this tracing is converted into a segmentation based on vessel boundary points that is not displayed here. For this result, $S_{vessel} = 0.94$, $F_{vessel} = 0.35$, $S_{background} = 0.94$, and $F_{background} = 0.002$. Based on these scores the latter algorithm successfully detected more vessel pixels, but with a higher incidence of false positives.

negative counts can be determined by using a distance tolerance, d. A true positive is a trace point in C that is within distance d of a trace point in G_{Trace}. A false positive is a trace point in C that has no corresponding points in G_{Trace} within distance d. A false negative would be a trace point in G_{Trace} that has no corresponding trace point in C within distance d. All others pixels not on the traces are true negatives.

4 Using Multiple Images to Improve Vessel Extraction

4.1 Mission

To accurately locate the regions in individual images that are blood vessels based on information obtained from equivalent regions in multiple images taken in the same patient sitting and to ensure that same regions in multiple images generate identical results.

4.2 Discussion

Before detecting longitudinal vessel change between two sets of images, it is essential that the vessels extracted from each individual image are as complete and accurate as possible. Unfortunately this is not always the case. When several images acquired from a single sitting are traced, (i.e. images that are temporally from the same set but are of different areas of the retina), it becomes apparent from examination of the results that there are significant differences. While the trace results are adequate for providing data necessary to perform registration, the results need to be made as accurate as possible prior to performing change detection. A method to determine the differences and for making the results equivalent needs to be examined.

These differences manifest themselves in two ways. First is the issue of tracing errors. These include missing traces (i.e. false negatives) and "phantom traces" (i.e. false positives). Vessels segments that are traced in one image are often missing in others. Both of these kinds of difference can be readily seen in Figure 11. Second is the issue of errors in the vessel boundaries. Although the vessels may be present in two images, their boundaries are not necessarily in agreement. Thus in order to detect longitudinal change, it first becomes necessary to make each set of images on the time-line as accurate as possible. To do this, it is necessary to detect disagreement in the extracted vessels, both in terms of

existence and in terms of boundary location, and to reconcile this disagreement.

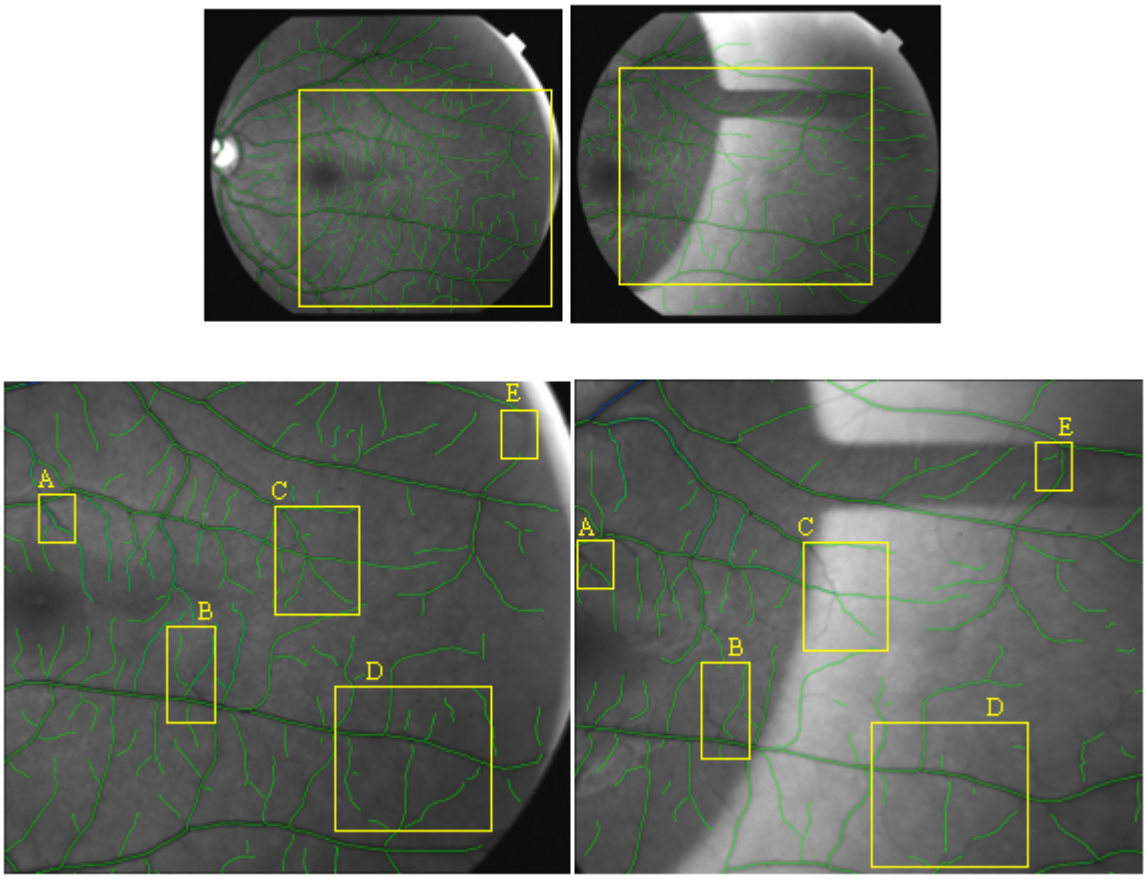

Figure 11: The top two images are images acquired in a single visit. The bottom images are close-ups of the boxed regions in the top images. Each of the rectangles in the lower images demonstrate how there is some variability in vessels that are detected. Note how the regions labeled A, B, C, and D in the left image have traces that are not present in the right image and how regions B, and E in the right image have traces that are missing in the left image. Clearly combining both sets of results would yield a more complete tracing but care needs to exercised so as not to increase the number of false positives.

4.3 A Proposed Methodology

Three steps are needed to generate more complete and accurate trace results. The first step is to run tracing on each individual image. This will be accomplished using Can's modified tracing algorithm discussed in Section 2. Using the features extracted during the

tracing of the images, the second step is to register the images using the Dual Bootstrap ICP Registration Algorithm. This results in a set of transformations that allow us to identify all corresponding points in all images. Armed with these transformations, the final step is to run tracing a second time, but this time simultaneously running tracing on all images and merging the results at each step of tracing for all corresponding points. At each step of tracing, each image that contains the area that is currently being traced will be consulted. If the results for the individual images are different, for instance, x images find no vessel and y images do, a resolution scheme will be used to decide how to precede. The resolution schemes used will vary. They may include any or all of the following:

- a voting strategy;
- a combining or averaging of the image intensities prior to attempting vessel detection;
- a combining or averaging of the template responses during vessel boundary detection.

During each of the above methods, individual image results may be weighted by a confidence score. This confidence score may be based on global or local contrast measures, vessel boundary strength, or prior success of previous iterations. For instance, it may be desirable to weight more favorably the input for regions of images with greater local contrast. However resolved, the final results will be reflected simultaneously on all images containing this region. This method proceeds until all regions on all images are traced with each image containing identical tracing results.

5 Vessel Change Detection

5.1 Mission

To accurately detect and quantify degree of blood vessel change in longitudinal sets of images.

5.2 Discussion

The ability to detect vessel change over time is a challenging task. This is due to the differences in images caused by the unavoidable circumstance that the distance and the angle from the camera to the retina varies as each image is acquired. The first circumstance causes the distance between any two corresponding points in different images to differ only in scale. The second circumstance causes projective differences that cause any two corresponding points in different images to differ in varying degrees. However both these circumstances can be compensated for through the process known as registration. By registering images with each other, transformations are created that can be used to convert images into a common scale-space. This ensures that the distance and direction between the any two corresponding points in the converted images are equal. Without this ability to transform images into a common scale space, it would not be possible to compare distances and hence detect changes in width. Fortunately, the technology now exists with the DBICP Registration algorithm to reliably and accurately register almost any set of fundus images that contain detectable features.

5.3 Previous Methods

The work of Giansanti et al [15] attempts to evaluate retinal changes by measuring the diameter and paths of vessels, determining the tortuosity of vessels and calculating the positions and angles of vessel crossings. However their methodology is only aimed at

measuring these indexes in a single image and provides no methodology for registration of results and for detection of change between images. Another shortcoming in their method is their tracing algorithm. It is limited because it only traces vessels starting from the optic nerve. However, not all vessels in an image field are connected to the vessels originating from the optic nerve and images often do not have the optic nerve.

Berger et al [2] detect vessel change by providing a user with a pair of registered images. Using these registered image pairs, they then suggest two methods. In the first method, these image pairs are made into slides that can either be superimposed or viewed with a stereo viewer. In the second method, which they term "alternation flicker", two registered images are viewed on a computer monitor in rapid succession at a rate of .3 to 10 Hz. Both methods are error prone in that they rely upon a user to detect the change. Additionally, their registration algorithm uses a "custom developed, nonrigid polynomial warping" algorithm that requires a user to manually pick six corresponding points between the two images. As such, their registration system is not as accurate or as robust as the DBICP algorithm which I will be using.

5.4 Proposed Methodology

In order to build an automated system with which to detect changes in vessel widths, three requirements must be met. First, there must be a method of finding the vessels in an image. Second, there needs to be a method of identifying corresponding vessel pieces. Third, there must be a way to measure and identify the changes in width. These requirements are depicted in Figure 12 and will be discussed in the following sections.

5.4.1 Vessel Detection

The first requirement in detection of vessel change is to be able to accurately and consistently identify vessels in images. This includes the accurate location and continuous

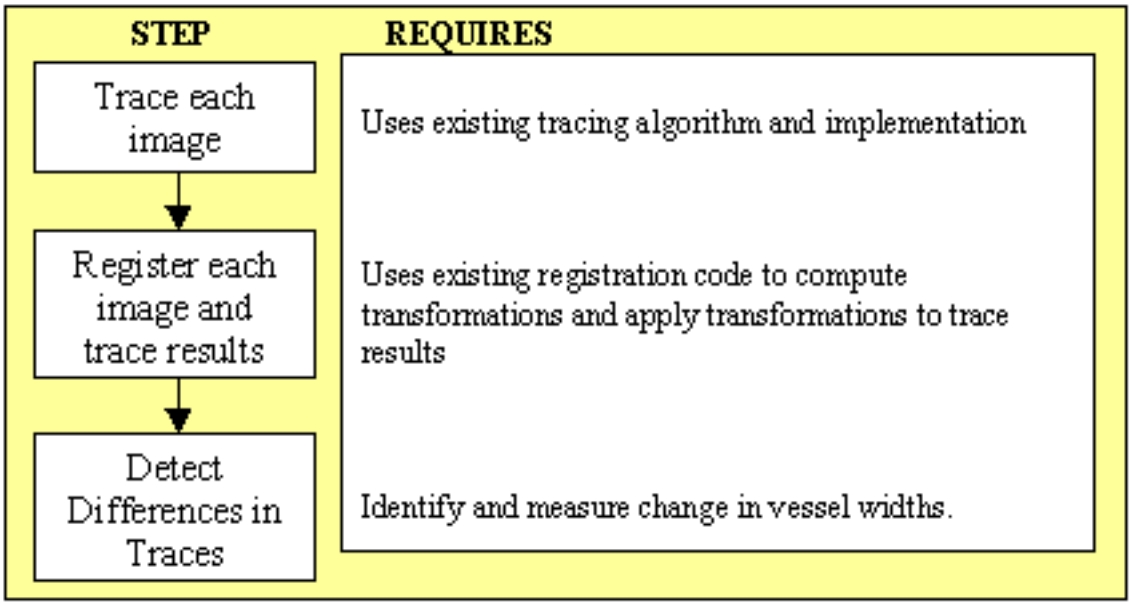

Figure 12: Showing the steps needed and the requirements to be able to use detected differences to construct better trace results from multiple images acquired in a single sitting.

definition of the vessel boundaries. To accomplish this, I plan to use the methods discussed in Sections 2 and 4

5.4.2 Determining Corresponding Vessels

The second requirement is the critical step in this process. It is the ability to accurately register images. By using the DBICP registration algorithm, it becomes possible to accurately align images. Once aligned, all the images are in a common scale-space and only then is it possible to accurately determine corresponding vessel cross sections and compare their widths. However, finding corresponding vessels is a challenge, particularly in the case where a vessel has disappeared or has failed to be detected in one image but is present in another. Thus a search strategy that recognizes this possibility needs to be considered when determining corresponding vessels.

One such search strategy would be as follows: If a vessel detected in one image coincides with a vessel in another image, then the vessels in both images should be accepted as being the same. If a vessel is within some distance δ of another vessel, then an additional constraint needs to be applied to ensure they are the same vessel. One such constraint could be to check if they are generally progressing in the same direction for a determined local area. If neither of these conditions are met, then it should be assumed that the vessel has either escaped detection or no longer exists. In this case, the image that is missing the vessel needs to be once again consulted to determine if the vessel is no longer there or if it was a failed detection. Vessel extraction needs to be attempted in the region of the missing vessel with different parameter settings (such as lower local thresholds). Only after a second attempt to detect the vessel should the vessel be considered as being no longer present.

5.4.3 Measuring Vessel Width

Many different techniques have been used to estimate vessel diameters, all of which are predicated upon the idea of measuring a vessel perpendicular to its local longitudinal orientation. Each method defines a cross section and defines the vessel boundaries on the cross section between which the width is measured. One such method, showm in Figure 13 is to determine vessel endpoints at the locations where the intensity is halfway between the maximum and minimum intensities of the cross sectional profile [9]. This is referred to as half-height intensity change. Others have used parameters of a parabola fit to a cross sectional intensity profile to determine the width [28] while others have fit gaussian curves to the profile and define the boundaries at some distance from the gaussian's mean [46, 13] as shown in Figure 13.

Most recently, Gang et al [12] have proposed the use of amplitude modified 2nd order gaussian filters to measure vessel width. They adaptively estimate the value for σ by using the following equation:

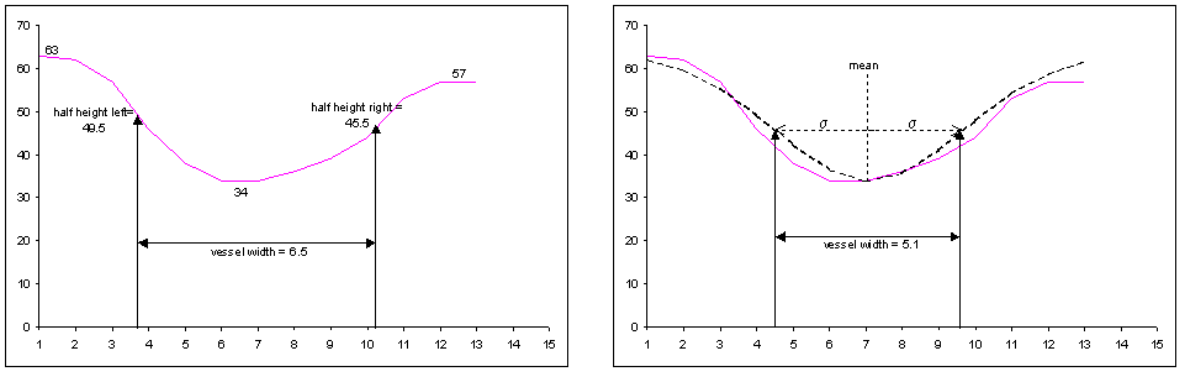

Figure 13: Examples of two methods of measuring width. The method on the left illustrates half height intensity change where the boundaries are the points on the intensity profile that are half way between the maximum and minimum intensity values. The method on the right shows a gaussian curve fit to the intensity profile where the boundaries are the points located at some distance from the mean, in this case 1 σ. Note how both methods generate different answers for the same intensity profile.

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma^3}(x^2 - \sigma^2)e^{-x^2/2\sigma^2}. \tag{13}$$

They demonstrate that such a filter will have a peak response for a vessel of a particular diameter when the correct value for σ is used. They assert that there is a linear relation between vessel diameter and σ defined by

$d = 2.03\sigma + .99$

where σ is set to the value used in the filter that generated the maximum response.

All of these methods suffer in that they arbitrarily define the boundaries based on thresholds (eg. σ) or estimates (eg. the curves fit to the profile or the levels used for half heights) and are affected by angular discretization errors.

To avoid these types of errors, I propose to measure widths directly as the distance between two points. These two points will be defined as points located on smooth curves representing the detected boundaries that are perpendicular to the vessel orientation and are located. By considering width measurements along corresponding points on corresponding vessels, widths can be compared to identify differences. If these differences

exceed the "normal" maximum expected difference to account for vessel change caused by the cardiac cycle of 4.8% [7], these vessel portions would be flagged as sites where differences have been observed.

In displaying the detected differences, the object is to draw the attention of the physician to the regions that appear to have the most change. As such, these results can be displayed in the form of a vessel difference map. This difference map would be a color representation of the vessels where the color of a vessel segment would be indicative of the amount of detected change. Regions that have the most significant colors can then be examined more closely by the physician.

5.5 Validating Vessel Widths

The accuracy of vessel width change detection is affected by the accuracy of the width measurement which in turn is affected by the accuracy of the boundaries. Thus demonstrating the validity of the the width measurement is a two step process. First the boundary detection needs to be proven as repeatable and then the width calculation needs to be proven to be repeatable.

To do this, I propose to detect vessels and boundaries from multiple images acquired from a single patient in a single sitting. I then will register the images and compare similar boundaries. The next step is to determine widths at selected sites in each image and compare widths. The only difference in width should be as a result of the cardiac cycle which has be shown to account for up to a 4.8% change in vessel width. Any other detected change can be attributed to the methodology used in measuring the width.

6 Specific Expected Contributions

In addition to providing algorithms for a diagnostic tool for physicians, several core technical innovations are expected from the project.

1. New model for detecting retina vessel boundaries.
2. Method for detecting width change in blood vessels between two images.
3. Method for improving completeness/accuracy of vessel detection using information from multiple images.
4. Methods for establishing fundus image segmentation probablistic gold standards.
5. Measures/metrics for comparing an algorithm's vessel segmentation with the gold standard or with another algorithm.

7 Schedule and Milestones

Milestone	**Expected Date of Completion**
1. Accurate Vessel Boundaries	November 2002
2. Multi-image Tracing	January 2003
3. Change Detection	March 2003
4. Gold Standard/Validation Method	May 2003
5. Thesis Writing	June 2003
6. Report to West Point	June 27, 2003
7. Thesis Defense	????

8 Current Assessment of Progress

Below is a current assessment of my progress in each area of my research.

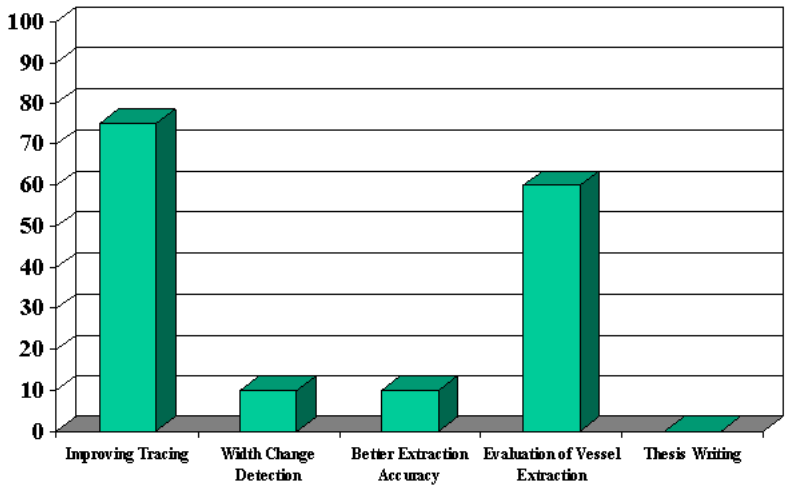

Publication: Jan-May 2002, authored book chapter [10] that described several retina vessel detection models and algorithms and provided examples of their applications. The abstract follows:

Quantitative morphometry of the retinal vasculature is of widespread interest, directly for ophthalmology, and indirectly for other diseases involving structural and/or functional changes of the body vasculature. Key points such as bifurcations and crossovers are of special interest to developmental biologists and clinicians examining conditions such as hypertension and diabetes. Segmentation/tracings of the retinal vasculature and the key points are also important as spatial landmarks for image registration. Image registration in turn has direct applications to change detection, mosaic synthesis, real-time tracking, and real-time spatial referencing. Change detection is important for supporting a variety of clinical trials, high-volume reading centers, and for large-scale screening applications.

The best-available algorithms for segmenting/tracing retinal vasculature are model based, and a variety of models are in use. Depending upon the intended application, different algorithmic and implementation choices can be made. This chapter will describe some of these models and algorithms and illustrate some of the implementation choices that need to be considered using a real-time algorithm as an example. Also described are methods for extracting key points, such as bifurcations and crossovers, and a discussion of how vessel morphometric data may be applied. Some methods are presented for generating ground truth or "gold standard" images as well as comparing these against computer generated results. Finally, some experimental analysis is presented for RPI-Trace and the ERPR landmark determination algorithms developed by our group.

References

[1] K. Akita and H. Kuga. A computer method of understanding ocular fundus images. *Pattern Recognition*, 15(6):431–443, 1982.

[2] J. W. Berger, T. R. Patel, D. S. Shin, J. R. Piltz, and R. A. Stone. Computerized stereo chronoscopy and alternation flicker to detect optic nerve head contour change. *Ophthalmology*, 107(7):1316–20, 2000.

[3] H. H. Brinchmann-Hansen, O. Theoretical relations between light streak characteristics and optical properties of retinal vessels. *Acta Ophthalmologica*, pages 33–37, 1986.

[4] A. Can, H. Shen, J. N. Turner, H. L. Tanenbaum, and B. Roysam. Rapid automated tracing and feature extraction from live high-resolution retinal fundus images using direct exploratory algorithms. *IEEE Trans. on Info. Tech. for Biomedicine*, 3(2):125–138, 1999.

[5] N. Chapman, N. Witt, X. Gao, A. Bharath, A. Stanton, S. Thom, and A. Hughes. Computer algorithms for the automated measurement of retinal arteriolar diameters. *British Journal of Ophthalmology*, 85:74–79, 2001.

[6] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson, and M. Goldbaum. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Transactions on Medical Imaging*, 8(3):263–269, 1989.

[7] H. Chen, V. Patel, J. Wiek, S. Rassam, and E. Kohner. Vessel diameter changes during the cardiac cycle. *Eye*, 8:97–103, 1994.

[8] O. Chutatape, L. Zheng, and S. Krishman. Retinal blood vessel detection and tracking by matched gaussian and kalman filters. In *Proceeding of IEEE Int Conf Engineering in Medicine and Biology Society*, volume 20, pages 3144–3149, 1998.

[9] F. Delori, K. Fitch, G. Feke, D. Deupree, and J. Weiter. Evaluation of micrometric and microdensitometric methods for measuring the width of retinal vessel images on fundus photographs. *Graefes Archive for Clinical and Experimental Ophthalmology*, 226:393–399, 1988. Manual methodology for measuing vessel widths.

[10] K. Fritzsche, A. Can, H. Shen, C. Tsai, J. Turner, H. Tanenbuam, C. Stewart, and B. Roysam. Automated model based segmentation, tracing and analysis of retinal vasculature from digital fundus images. In J. S. Suri and S. Laxminarayan, editors, *State-of-The-Art Angiography, Applications and Plaque Imaging Using MR, CT, Ultrasound and X-rays*. Academic Press, 2002.

[11] P. Fua and Y. G. Leclerc. Model driven edge detection. *Machine Vision and Applications*, 3:45–56, 1990.

[12] L. Gang, O. Chutatape, and S. Krishnan. Detection and measurement of retinal vessels in fundus images using amplitude modified second-order gaussian filter. *IEEE Transactions on Biomedical Engineering*, 49(2):168–172, 2002. Vessels found via matched filter and widths calculated with amplitude modified second order gaussian equation.

[13] X. Gao, A. Bharath, A. Stanton, A. Hughes, N. Chapman, and S. Thom. Measurement of vessel diameters on retinal images for cardiovascular studies. In *Proceedings of Medical Image Understanding and Analysis*, 2001.

[14] X. Gao, A. Bharath, A. Stanton, A. Hughes, N. Chapman, and S. Thom. A method of vessel tracking for vessel diameter measurement on retinal images. In *Proceedings IEEE International Conference on Image Processing 2001*, pages 881–884, 2001.

[15] R. Giansanti, P. Fumelli, G. Passerini, and P. Zingaretti. Imaging system for retinal change evaluation. *Sixth International Conference on Image Processing and Its Applications*, 2:530–534, 1997.

[16] M. Hammer, S. Leistritz, L. Leistritz, D. Schweitzer, and E. Thamm. Monte-carlo simulation of retinal vessel profiles for the interpretation of in viva oxymetric measurements by imaging fundus reflectometry. In *Proceedings SPIE Medical Apps of Lasers in Derm, Ophth, Dent,& Endo*, volume 3192, pages 211–218, 1997.

[17] A. Hoover, V. Kouznetsova, and M. Goldbaum. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Transactions on Medical Imaging*, 19(3):203–210, 2000.

[18] A. Houben, M. Canoy, H. Paling, P. Derhaag, and P. de Leeuw. Quantitative analysis of retinal vascular changes in essential and renovascular hypertension. *Journal of Hypertension*, 13:1729–1733, 1995.

[19] J. Hu, R. Kahsi, D. Lopresti, G. Nagy, and G. Wilfong. Why table ground-truthing is hard. In *Proceedings of the Sixth International Conference on Document Analysis and Recogniton*, pages 129–133, 2001.

[20] J. Jagoe, C. Blauth, P. Smith, J. Arnold, K. Taylor, and W. R. Quantification of retinal damage done during cardiopulmonary bypass: Comparison of computer and human assessmement. *IEE Proceedings Communications, Speech and Vision*, 137(3):170–175, 1990.

[21] P. Jasiobedzki, D. McLeod, and C. Taylor. Detection of non-perfused zones in retinal images. *Computer-Based Medical Systems: Fourth Annual IEEE Symposium*, pages 162–169, 1991.

[22] P. Jasiobedzki, C. Taylor, and J. Brunt. Automated analysis of retinal images. *Image, Vision and Computing*, 3(11):139–144, 1993.

[23] P. Jasiobedzki, C. Williams, and F. Lu. Detecting and reconstructing vascular trees in retinal images. In M. Loew, editor, *SPIE Conference on Image Processing*, volume 2167, pages 815–825, 1994.

[24] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[25] B. Kochner, D. Schuhmann, M. Michaelis, G. Mann, and K. Englmeier. Course tracking and contour extraction of retinal vessels from color fundus photographs: Most efficient use of steerable filters for model-based image analysis. In *SPIE Conference on Image Processing*, volume 3338, pages 755–761, 1998.

[26] H. Li and O. Chutatape. Fundus image features extraction. In J. Enderle, editor, *Proc. IEEE Int. Conf. Engineering in Medicine and Biology*, volume 4, pages 3071 – 3073, Chicago, IL, USA, July 2000. Intelligent Machines Lab., Nanyang Technol. Univ., Singapore.

[27] H. Mayer, I. Laptev, A. Baumgartner, and C. Steger. Automatic road extraction based on multiscale modeling. *International Archives of Photogrammetry and Remote Sensing*, 32:47–56, 1997.

[28] F. Miles and A. Nutall. Matched filter estimation of serial blood vessel diameters from video images. *tmi*, 12(2):147–152, 1993.

[29] W. Neuenschwander, P. Fua, L. Iverson, G. Szekely, and O. Kubler. Ziplock snakes. *International Journal of Computer Vision*, 25(3):191–201, 1997.

[30] W. Niessen, C. Bouma, K. Vincken, and M. Viergever. Error metrics for quantitative evaluation of medical image segmentation. In R. Klette, H. Stiehl, M. Viergever,

and K. Vincken, editors, *Performance Characterization in Computer Vision*, pages 275–284. Kluwer Academic Publishers, The Netherlands, 2000.

[31] L. Pedersen, M. Grunkin, B. Ersbll, K. Madsen, M. Larsen, N. Christoffersen, and U. Skands. Quantitative measurement of changes in retinal vessel diameter in ocular fundus images. *Pattern Recognition*, 21:1215–1223, 2000.

[32] A. Pinz, S. Bernogger, P. Datlinger, and A. Kruger. Mapping the human retina. *IEEE Transactions on Medical Imaging*, 17(4):606–620, Aug 1998.

[33] D. Roberts. Analysis of vessel absorption profiles in retinal oximetry. *medphys*, 14(1):124–139, 1987. Vessel profile discussion. Uses 4 parameter curve fitting procrdure. Addresses and describes central reflex.

[34] C. Sinthanayothin, J. Boyce, H. Cook, and T. Williamson. Automated localisation of the optic disc, fovea, and retinal blood vessels from digital colour fundus images. *The British Journal of Ophthalmology*, 83(3):902–10, Aug 1999.

[35] N. Solouma, A. Youssef, Y. Badr, and Y. Kadah. Real-time retinal tracking for laser treatment planning and administration. In *SPIE Conference Medical Imaging:Image Processing*, volume 4322, pages 1311–1321, 2001.

[36] A. Stanton, B. Wasan, A. Cerutti, S. Ford, R. Marsh, P. Sever, S. Thom, and A. Hughes. Vascular network changes in the retina with age and hypertension. *Journal of Hypertension*, 13(12):1724–1728, 1995.

[37] C. Stewart, C.-L. Tsai, and B. Roysam. The bootstrap iterative closest point algorithm with application to retinal image registration. *International Journal of Computer Vision*, submitted 2002.

[38] I. Suzuma, Y. Hata, allen Clermont, F. Pokras, S. Rook, K. Suzuma, E. Feener, and L. Aiello. Cyclic stretch and hypertension induce retinal expression of vascular

endothelial growth factor and vascular endothelial growth factor receptor-2: potential mechanisms for exacerbation of diabetic retinopathy by hypertension. *Diabetes*, 50:444–454, Feb 2001. Good source for other sources/citations.

[39] W. Tan, Y. Wang, and S. Lee. Retinal blood vessel detection using frequency analysis and local-mean-interpolation filters. In *SPIE Conference Medical Imaging:Image Processing*, volume 4322, pages 1373–1384, 2001.

[40] G. Tascini, G. Passerini, P. Puliti, and P. Zingaretti. Retina vascular network recognition. In *SPIE Conference on Image Processing*, volume 1898, pages 322–329, 1993.

[41] Y. Wang and S. Lee. A fast method for automated detection of blood vessels in retinal images. In M. Fargues and R. Hippenstiel, editors, *Signals, Systems & Computers, 1997*, volume 2, pages 1700 – 1704, Nov. 1997.

[42] T. Wong, R. Klein, A. Sharrett, B. Duncan, D. Couper, J. Tielsch, B. Klein, and L. Hubbard. Retinal arteriolar narrowing and risk of coronary heart disease in men and women the atherosclerosis risk in communities study. *Journal of American Medical Association*, 287(9):1153–1159, March 6 2002.

[43] T. Y. Wong, R. Klein, A. R. Sharrett, M. I. Schmidt, J. S. Pankow, D. J. Couper, B. E. K. Klein, L. D. Hubbard, and B. B. Duncan. Retinal arteriolar narrowing and risk of diabetes mellitus in middle-aged persons. *Journal of American Medical Association*, 287(19):2528–2533, May 15 2002.

[44] F. Zana and J.-C. Klein. Robust segmentation of vessels from retinal angiography. In *Proceeding International Conference Digital Signal Processing*, pages 1087–1090, 1997.

[45] F. Zana and J.-C. Klein. Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE Transactions on Image Processing*, 10(7):1010–1019, 2001.

[46] L. Zhou, M. Rzeszotarski, L. Singerman, and J. Chokreff. The detection and quantification of retinopathy using digital angiograms. *IEEE Transactions on Medical Imaging*, pages 619 – 626, 1994.