FORTH-ICS / TR-416 May 2011

Retinal Images Analyzer †

Polykarpos Karamaounas and Georgios Manikis
and Xenophon Zabulis

Polykarpos Karamaounas and Georgios Manikis
and Xenophon Zabulis

Institute of Computer Science
Foundation for Research and Technology — Hellas (FORTH)
N. Plastira 100, Vassilika Vouton
Heraklion, Crete, 700 13 Greece

Web: `http://www.ics.forth.gr/`
E-mail: {karam | gmanikis | zabulis}@ics.forth.gr
Tel: +30 2810 391600, Fax: +30 2810 391601

Technical Report FORTH-ICS / TR-416— May 2011

Abstract

A software application is presented, for the detection and measurement of retinal vessels in fundoscopy images. This software is comprised by a measurement module and a Graphical User Interface (GUI) module. The measurement module employs image processing methods and operates in batch mode, to segment retinal vessels in the input image. In addition, the measurement module estimates the diameters of the imaged vessels along their spatial extent in the image, based on the above segmentation. The user interface module provides methods for the automatic and interactive measurement of vessel diameters at points and regions of interest. Moreover, it provides methods for editing the vessel representation, in order to recover from segmentation errors. These two modules are integrated into a medical application. In turn, this application increases the automation of a number of vessel measurement protocols, which are of interest to medical professionals.

1 Introduction

The retina provides an open and accessible window for studying the microcirculation in the human body. Retinal vessels can be easily visualized with non-invasive techniques providing information corresponding to the anatomical and physiological properties of cerebral and coronary microcirculation [1]. According to recent large epidemiological studies, hypertensive retinopathy signs are associated with blood pressure and specific target organ damages, indicating that evaluation of retinal blood vessels could serve as a cardiovascular risk predictor of great value.

Assessment of retinal vessels in hypertension is made possible by fundoscopy, which includes an overview of the retina and vitreous of the eye. Fundoscopy can be performed with a direct or indirect ophthalmoscope or with a slit lamp which, in addition, allows the three-dimensional imaging of the seabed. All three methods require prior medical mydriasis of the eye, specially trained examiners, are subjective (as images are not recorded) and do not permit the assessment of more sophisticated measurements of retinal vessels, such as the diameter of the vessels. In recent years, the construction of fundus cameras (mydriatic or non-mydriatic) facilitated the representation of retinal blood vessels, providing other possibilities for the diagnosis and classification of retinopathy. Non-mydriatic cameras exhibit the advantage of operating without a pharmacy dilation of the pupil, thereby significantly reducing the time of the examination and being better tolerated by the patient. Furthermore, the acquired images can be analyzed by software which measures the diameter of blood vessels, the application of specific measurement protocols, and are very reliable measures of hypertensive retinopathy even in the early stages.

Measurement of diameters by operator identification of vessel edges is a time consuming process and potentially subject to error and observer bias [12]. To overcome these issues attempts were made to be developed semi-automatic software that cast measurements faster and more repeatable, thus having more reliable results compared to manual measurements [9].

The contributions presented in this technical report comprise two software modules, which are integrated in a single application. First is the detection and measurement of vessels in retinal images, an image analysis task. Second, is the provision of the ability to validate, edit and represent this information in a way that (a) facilitates the interactive measurement of vessel and vessel segments of interest and (b) the extraction of statistical data of vessel segments within spatial regions dictated by pertinent medical protocols. While the first module is fully automatic, the functionalities of second module are provided through a Graphical User Interface (GUI). This GUI incorporates methods that increase the automation of the measurement process and render it user-friendly. Such methods include automatic segmentation and measurement of a vessel-segment of interest, the computation of statistics across user-defined regions of medical interest, the interactive correction and editing of segmentation results, the characterization of vessel segments with attributes (i.e. vain, artery) and other.

Given a fundoscopy image, our approach to the problem is the following:

1. Compute a measure that represents the likelihood of the pixel to be imaging a vessel. Result is a floating-point image A.
2. Segment vessels from image A. Result is binary image B.
3. Find the skeleton of image B. Result is image S.
4. Measure vessel width at the skeleton points found in image S.
5. Avail a user interface that provides methods for the measurements of interest, both for individual points as well as across user-selected regions of interest (vessel segments, or retinal regions).

Fig. 1 and Fig. 2 provide a brief overview of system functionalities.

The remainder of this report is organized as follows. In Section 2 related work is reviewed. In Section 3, we present our approach the vessel segmentation of fundoscopy images and the skeletonized representation of the segmented vessels. In Section 4, the employed vessel measurement techniques, which are based on this skeletonized representation, are presented. In Section 5, we present the Graphical User interface of the application, which employs methods that automate specific measurements dictated by medical protocols and provide interactive measurements across user-selected regions of interest. Section 6, summarizes this work and provides directions for future work.

2 Relevant work

2.1 Medical measurements

To date, several semiautomatic method for the objective quantification of retinal microvascular characteristics, i.e. [13, 4]. The most widely used software was developed and validated in the Atherosclerosis Risk in Communities Study [5], while the same methodology was followed in [14].

2.2 Work in image processing for vessel extraction

Detecting and measuring spatial properties of the vessel network through automated techniques is a challenging task, but which can support diagnosis. From the perspective of the image processing, much effort has been paid for the segmentation of the retinal vasculature and the extraction of its anatomical data, such as width, length etc. A number of vessel detection and measurement techniques have been suggested in the literature and analytically described in [8]. The approaches can be classified in two broad categories. In the first, lie methods which segment vessels from the image as a first step to estimate retinal structure. On the second lie methods that track the vessel boundaries to extract the shape of the former.

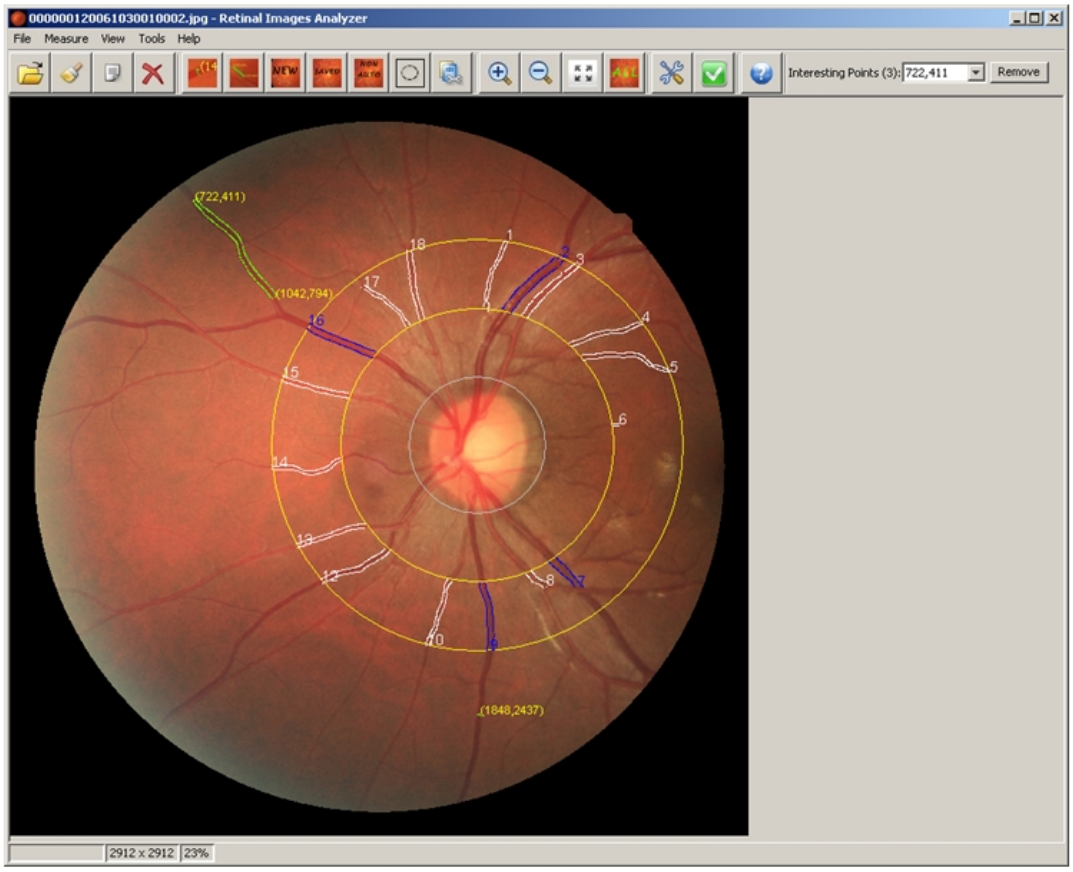

Figure 1: Main user interface UI panel. An image has been loaded and analyzed. A gray circle marks the estimate of the optical disk. The two concentric yellow circles around it, delimit a region of interest to the user. Inside this region vessels have been automatically detected, enumerated, and measured; vessels 7 and 8 are delimited by the junction closest to the optical disk in accordance to user preferences. Using the UI the user has characterized vessels; arteries are intonated in the UI with white and veins with blue. Furthermore, the user has indicated a vessel-segment of interest at the upper left part of the image, which has been also traced and measured. All operations are either automatic or require minimal user input. On the upper part of the figure shown is a toolbar that provides the functionalities in a user-friendly manner. The system incorporates image analysis methods within the UI in order to provide high-level measurement tools, in a "point and click" fashion.

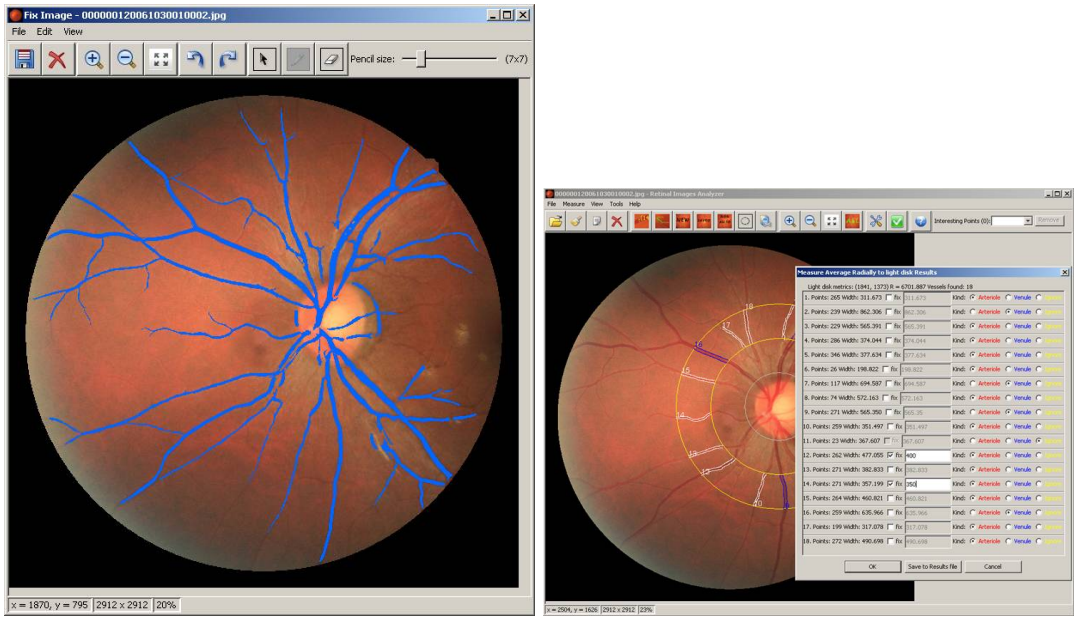

Figure 2: Secondary UI panels. Though the UI provides simple access to complex representations, the user is availed with full access to internal system representations. In this way, the user can correct potential shortcomings of image analysis and resolve ambiguities. In the left panel, the user can edit the segmentation result. At the right panel, the user can annotate an image with measurements on particular vessel-segments and export the results.

The entire pool of these techniques is generally divided into two groups. In the first group, pixel-processing based techniques such as adaptive local thresholding, global matched filtering and differential filtering fundamentally identify the vessel regions before exporting the available vasculature geometrical knowledge. More computationally complex techniques based on 2D wavelet transforms, supervised learning (K-nearest neighbors, Bayesian classifiers and neural networks) and active contours are also employed and belong to this category. The proposed work follows the pixel-based approach.

On the other hand, tracking methods focus on the identification of the vascular network, where the segmentation is not a prerequisite step for the identification. Given an initial position, the vessels detection is achieved by tracking their boundaries. Tracking methods were extended by the use of fuzzy C-means clustering. In such algorithms, the extraction of the vasculature requires no prior assumptions about the characteristics and the edge information of the vessels.

2.3 This work

In this work, a semi-local approach is employed to classify pixels as imaging vessels or not and the result of this segmentation is skeletonized in order to obtain a representation of the imaged vessels. As the segmentation result is not perfect, this skeleton is further processed in order to eliminate artifacts due to failures of segmentation. In addition, a user-interface is provided to semi-automatically correct residual errors.

3 Image processing

Retinal photographs were acquired with a NIDEK AFC-230/210 non-mydriatic digital fundus camera. The images acquired by this process are the input to this module, in the form of 2912×2912 pixel RGB color images in JPEG format. Given the input the image processing module automatically performs vessel detection through image segmentation, vessel skeletonization, and width measurement across the detected vessels.

3.1 Preprocessing

Initially, the acquired retinal images are transformed from Red-Green-Blue (RGB) to monochromatic, since further processing operates on the intensity values. Due to the fact that the vessels appear more contrasted in the green channel, the intensity of the green channel was adopted as the monochromatic image intensity.

Contrast limited adaptive histogram equalization (CLAHE) [16] is then applied, to increase the global image contrast and reveal as much as possible the structural information of the vasculature. The enhancement of the image contrast was implemented, employing the built-in function in Matlab "adapthisteq". The raw and the enhanced image (see Fig. 3) are presented below.

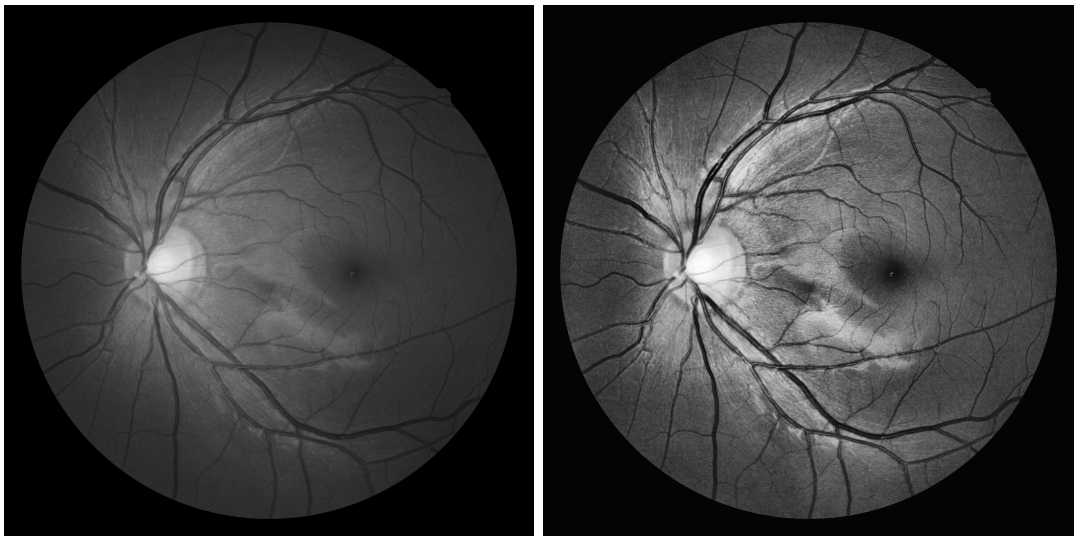

Figure 3: Adaptive histogram equalization. Original (left) and processed image (right).

3.2 Segmentation

Given the preprocessed image, the segmentation of blood vessels is performed in two steps:

1. vessel enhancement filtering, and
2. region-based identification of the vascular network.

The goal of the vessel enhancement filtering is to provide clearly defined vascular structures even for non-visible small vessels. The Hessian matrix of the input image is computed in order to quantify the curvature of the intensity profile as a ridge, using the metric at [3]. Note that the Hessian matrix describes the second order intensity variations around each point of the image.

Processing takes place as follows: initially the above metric is applied to the equalized image. The filtering result is then thresholded and cleared for a cyclic border that occurs in the periphery of the result. Fig. 4 illustrates the process.

3.3 Skeletonization

A post-processing step, required for the characterization of the morphological structure of the blood vessel's network, is the skeletonization of the binary image that is provided after the completeness of the segmentation procedure. Skeletonization algorithms generate one-pixel-wide skeletons by replacing the patterns that belong to class 1 (white) with line drawing representation of them. The morphological operation of skeletonization was implemented according to the built-in function "bwmorph" in Matlab [6]. The skeletonized image highlights relevant features of the vasculature

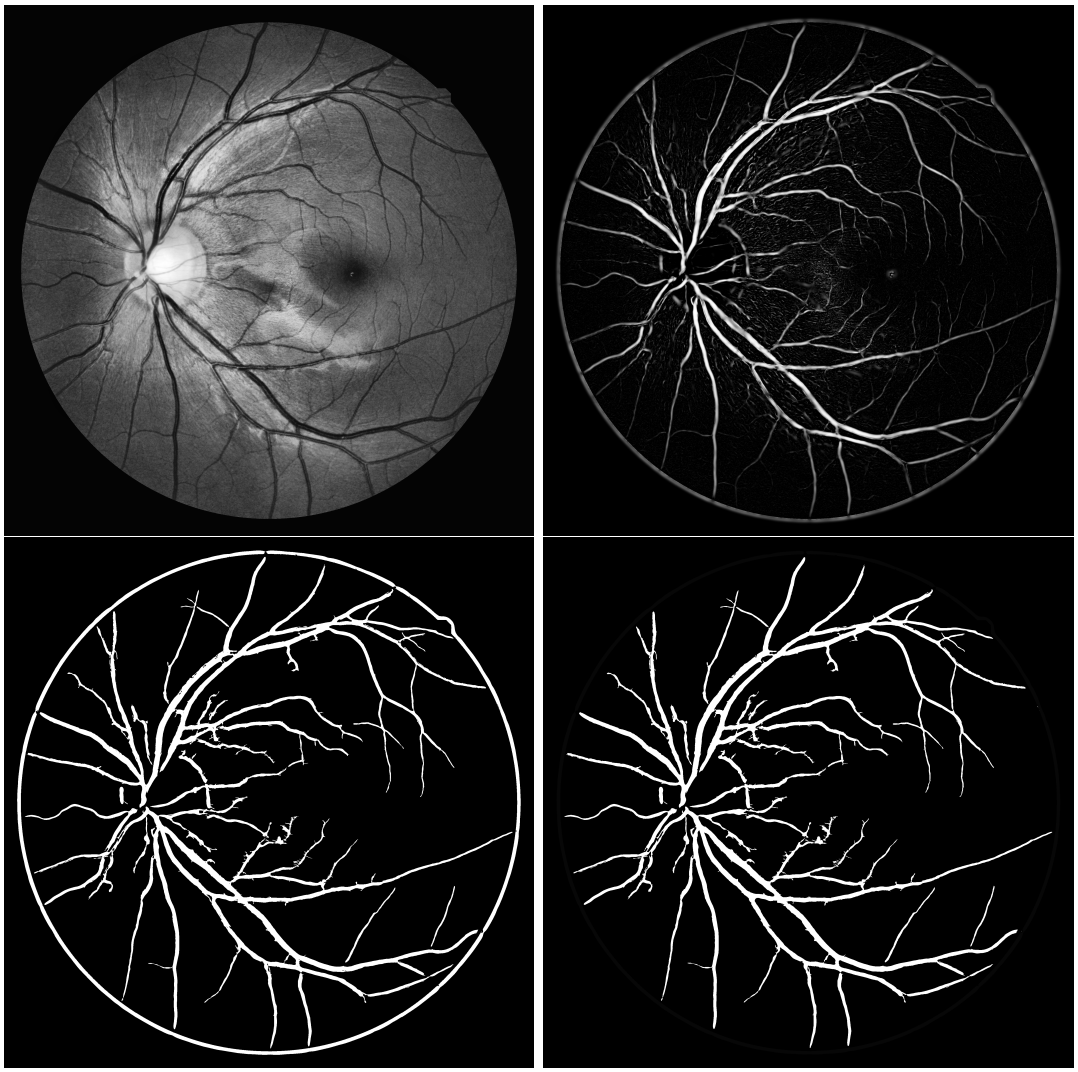

Figure 4: Retinal vessel segmentation. Top-left: input, equalized image. Top-right: "vessel-enhanced" image. Bottom-left shows the thresholding on the enhanced image and bottom-right the, cleared from border, segmentation result.

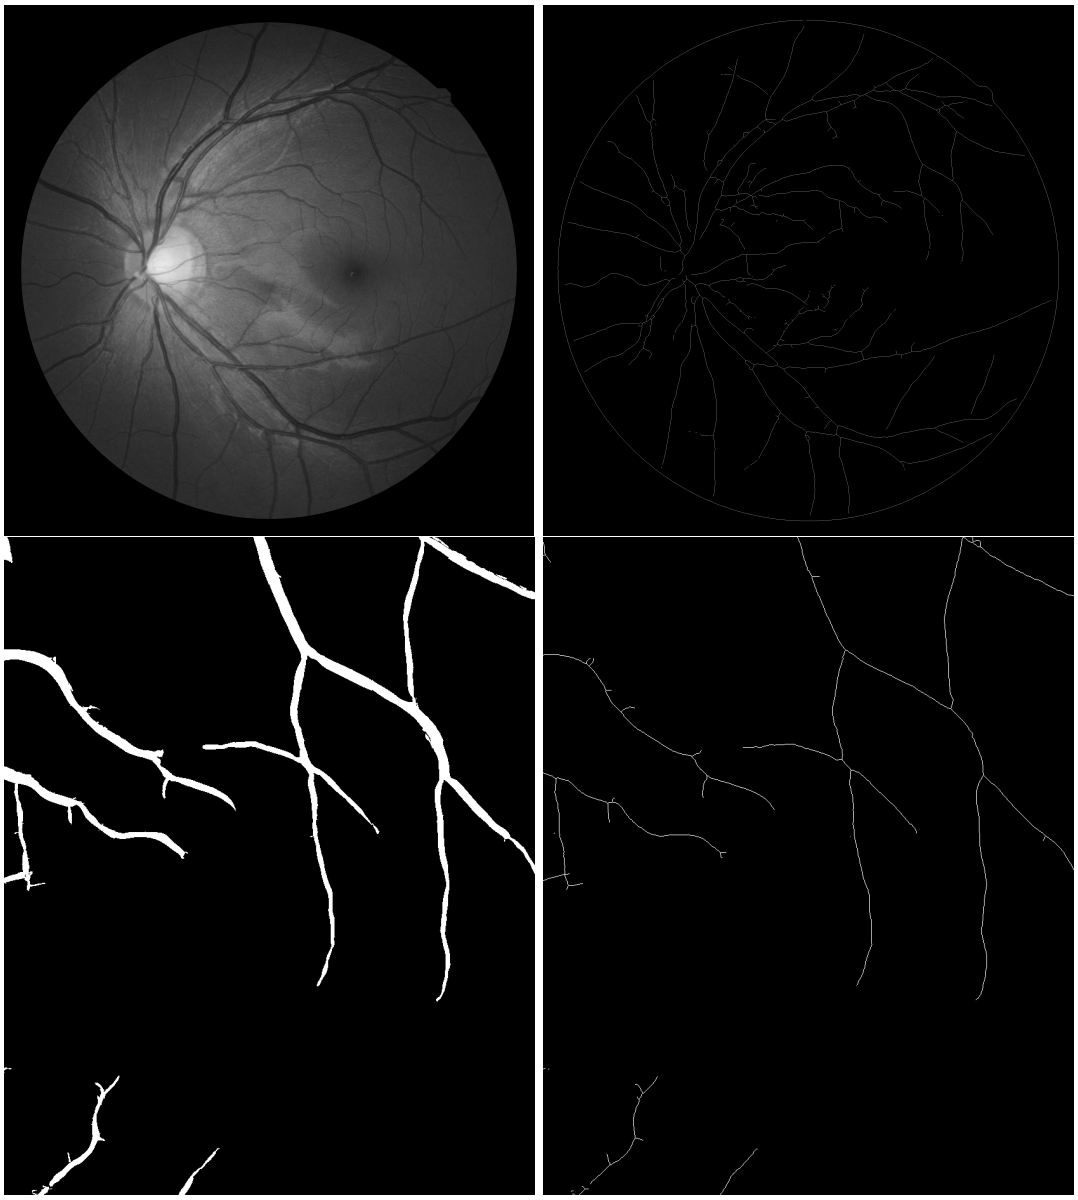

Figure 5: Skeletonization of original image. The top row shows the original image and the extracted skeleton. The bottom row illustrates skeletonization in a detail of the image: the left figure shows the segmentation and the right its skeletonization.

such as end-points, junction-points, and connection among the components. Below, we present the result of skeletonization (see Fig. 5).

3.4 Processing in parts

The retinal images from the clinicians had a resolution of 2900X2900 pixels. In general, such resolutions offer high quality image visualization but on the other hand,

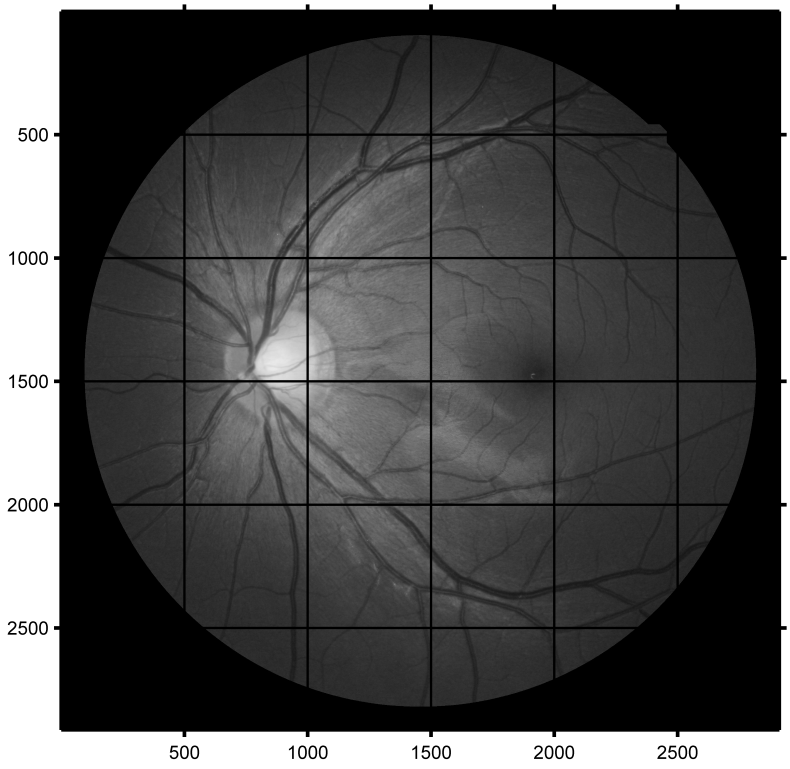

Figure 6: Block based image partitioning. In order to achieve the segmentation of high resolution images and cope with memory constraints, images are partitioned into blocks. These blocks are overlapping, in order to seamlessly perform semi-local operations such as image smoothing and differentiation.

their processing is a computationally demanding and time consuming approach. Concerning the necessity of providing as much accurate measures of the vascular network as possible, our approach focuses on the solution of the computational problems rather than downsizing the available images.

A common way to overcome the high-dimensionality issues is the image processing through distinct blocking. In distinct block processing, the image is divided into m-by-n blocks. These blocks, overlay the image matrix starting in the upper left corner, by adding some border pads when the block does not fit exactly over it. To perform distinct block operations, the "blockproc" built-in Matlab function was adopted. Furthermore, the function pads overlapping borders to each block so that local operations such as spatial filtering can be seamlessly performed. Finally, each distinct block is extracted from the image and passed to the segmentation algorithm. The "blockproc" function assembles the returned blocks to create the resulted image. Image Fig. 6 represent the block operation for a random processed image. Note that some borders are also added in each block.

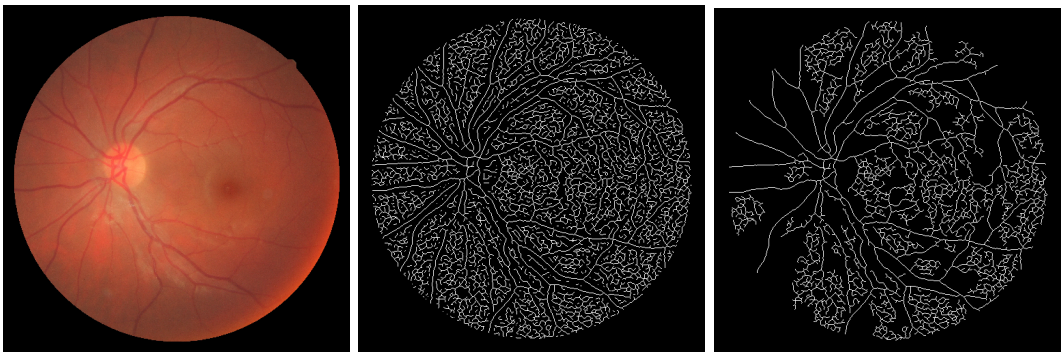

Figure 7: Elimination of spurious vessel segments, based on size. Left to right: original image, skeleton image, and filtered skeleton image.

4 Skeleton representation processing

After image segmentation and skeletonization, a number of operations is automatically performed. These processes operate on the segmented and skeletonized images to create an accurate representation of vessels and estimate quantities requited henceforth, such as the location of the optical disk in the image. The result of these processes is a priori computed and stored along with the image. It is loaded when the user is to perform measurements on an image. Since the user has the ability to modify the segmentation result the operations described in this section are re-executed for each the segmentation result is modified.

4.1 Size-based filtering

The first stage of pre-processing is to delete very small isolated skeleton segments as they typically correspond to noise artifacts. The minimum component size, segments below which will be eliminated, is an input parameter, whose default value is determined proportionally to the resolution of the image; for the case of a 2048×2048 image it was set to 200 pixels. A typical 8-connected components process is applied to identify isolated elements in the skeleton image. The output of the method is an updated skeleton image B. Is pseudocode is provided in the table below. Fig. 7 illustrates the operation.

The algorithm for this operation, is as follows.

For each pixel P of the skeletonized image that has not been parsed visited:

- Create a new connected component starting from pixel P and expand recursively blob for each one of the 8 adjacent pixels of P which is also a skeleton pixel and not already parsed.
- In case the extracted blob has length smaller than *threshold*, exclude it from the skeleton image by assigning a zero value to the corresponding pixels.

4.2 Prunning of small vessel branches

The segmentation results obtained from high resolution images tend to exhibit richer structure at vessel boundaries than corresponding results from low resolution images. In particular a number of intrusions and protrusions can be observed which are typically due to combinations of image noise and low image contrast. In low resolution images, such artefacts have a smaller spatial extent and do not produce additional branches in the extracted skeleton structure. However, in high resolution images such structures give rise to spurious skeleton branches (see in Fig. 8). The technique prunes "small" skeleton branches as they correspond typically these spurious branches.

At preprocessing such branches are pruned in order to facilitate more intuitive user interaction in the selection of branch segments. Given as input the skeleton image and the maximum branch length mL as an input parameter, a process prunes branches that have length smaller than mL and updates the skeleton image. The pseudocode for this operation is shown below. Fig. 8 illustrates the operation.

For each skeleton point (SP) which is also a junction point (adjacent points are more than 2):

- for each adjacent skeleton point (ASP) of junction point (SP):
  - initialize a path (P) with points SP and ASP and expand it recursively until:
    - ∗ end is reached (cannot be expanded more), or
    - ∗ another junction point is reached, or
    - ∗ the paths cardinality became greater than mL
  - in case path (P) reached an end (cannot be expanded more) and has length less than mL, this means that it is a "small" branch. Remove it from skeleton: remove all its points except the 1st one

The effect of this filtering provides a more accurate representation of vessels, particularly because vessel branching points convey important information in measurements (i.e. in Sec. 5 vessels are measured up to the first encountered branching points, according to the particular medical protocol). By eliminating the above artifacts the

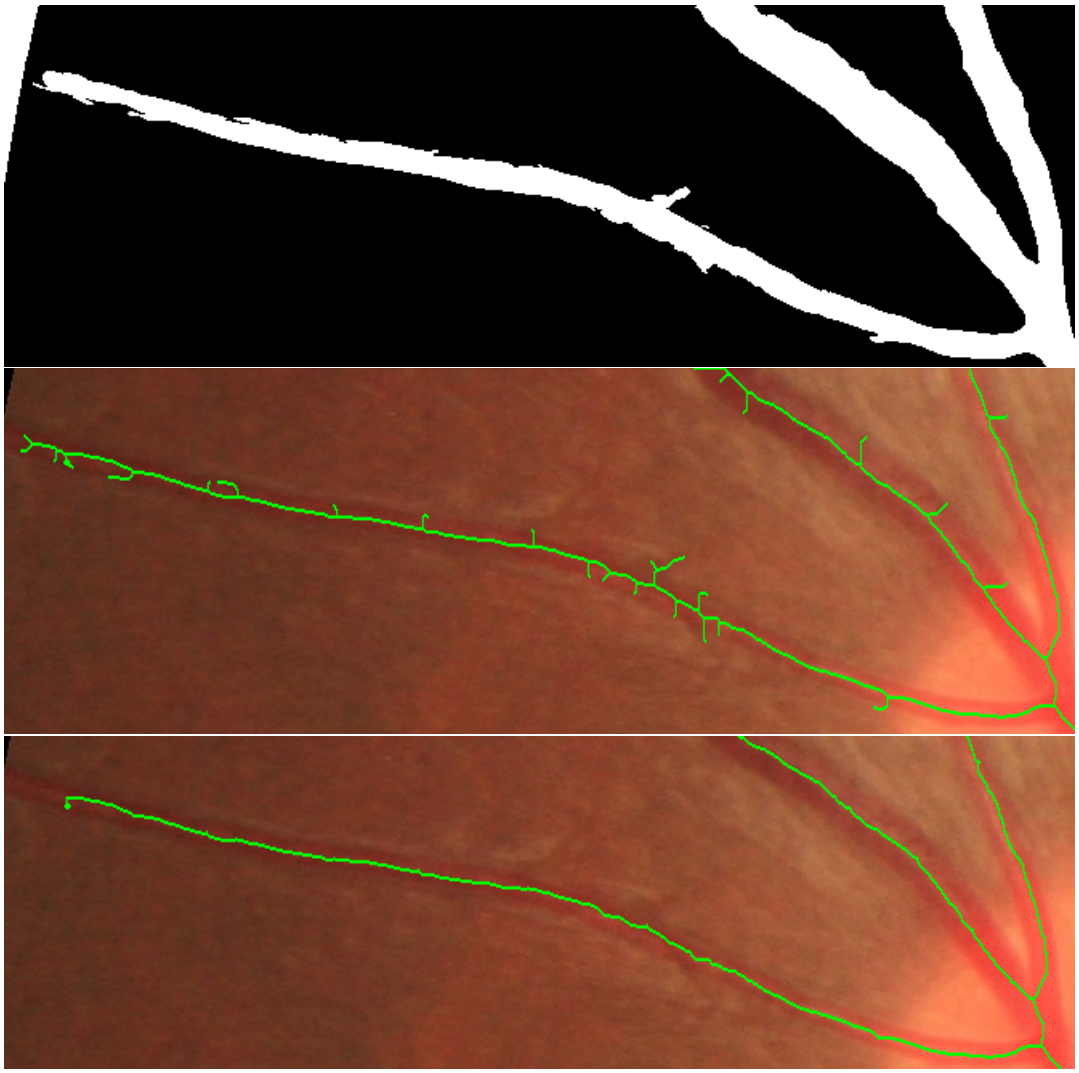

Figure 8: Pruning of spurious vessel branches based on size. Top to bottom: segmentation image, original image with initial skeleton superimposed, and original image with pruned skeleton superimposed.

representation is relieved from spurious branching points.

4.3 Vessel width estimation

A local algorithm is employed to estimate the vessel image width for vessels for which a skeleton point p is extracted. The algorithm operates on the skeleton and the segmentation images.

A hypothetical straight line segment T of length l is placed horizontally and centered upon p; length T is selected to be larger than any expected vessel width. Segment T is then consecutively rotated and for each rotation its intersections with the vessel's

boundaries are found. The rotation that minimizes this length is assumed to bring T in a posture that intersects the vessel perpendicularly to its local orientation axis. The distance of the two intersections for this rotation is the measurement of vessel width at point p.

To find the intersections of T with the vessel for a given rotation, we initiate our query from point p. We then visit pixels in image B in the direction of the rotated line segment, in both directions. The query, in each direction, ends when a non-vessel pixel (its value being 0) is found.

This computation is performed for all skeleton points and the results stored, thus being readily available when the image is loaded. In this way, the measurements are readily available when the image is inspected by the medical professional, thus increasing the interactivity of the application. When a user indicates a point (i.e. by clicking the mouse device) the closest skeleton point to it is retrieved and the corresponding measurement, for that skeleton point, is provided. Based on camera calibration, the measured length is converted from pixels to μm.

Finally notice that even if the skeleton point is inaccurately found, the rotation that minimizes the distance of the two intersection points width is still the same. Therefore, the proposed width estimation is invariant to small inaccuracies of the skeletonization algorithm. The pseudocode below described the operation and Fig. 9 illustrates the workings of the algorithm and its result. The input to this process is the skeletonized image S and the vessel segmentation image B. The output is the coordinates of each skeleton point and a floating point number representing the estimated width at this point, as saved in an output file.

Define a hypothetical segment T of length greater than the width of any vessel expected to be encountered (input parameter of the algorithm).
For each pixel p of the skeletonized image:

- Locate segment T horizontally centered at p
- For each rotation angle $0 - 180^o$ calculate the distance of the segment points (L, R) that intersect the vessel boundaries. To identify these points, start from point p and sequentially visit pixels along T in both directions until the vessel boundary is reached.
- The final vessel width (w) is the minimum distance found from the above rotation

Since measurements are performed in the digital image, corresponding width values are expressed in pixels. This is sufficient for most diagnostic purposes as medical professionals are interested in the variation of vessel width along time, rather than its metric value. Nevertheless, vessel width measured in pixels can be converted to metric values if the image is calibrated with reference object of known size. For this

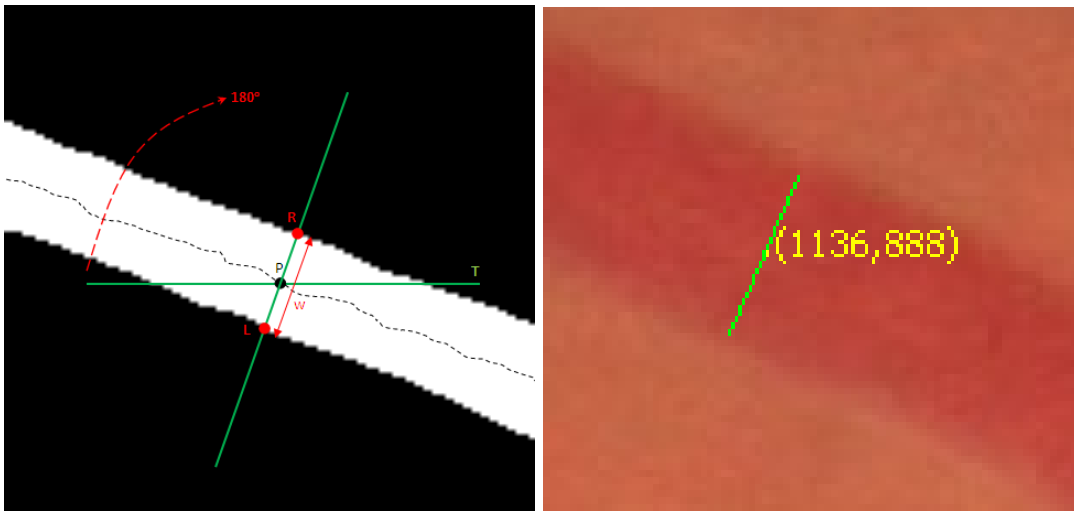

Figure 9: Illustration of vessel width measurement technique (see text).

purpose, the application provides the option of accounting for such calibration data and exporting vessel widths in the corresponding units.

4.4 Optic disk detection

A functionality is provided that automatically detects the optic disk in the acquired image and provides an estimate of its size. In this way, the application of medical protocols that are based on measurements around this disk can be automated.

The optic disk is detected as the brightest, size dominant blob in the image, as follows. Initially, several Mean Shift [2] processes are seeded equidistantly across the image. These processes use a circular kernel to converge to the local intensity maximum. Results leading to a low intensity value are disregarded. The remaining results are topologically clustered based on distance. The centroid C of the cluster with the largest cardinality is selected as the center of the optical disk. The process is described in more detail in [15] where is, inversely, applied to detect the size-dominant darkest blob of a circular image.

To estimate the size of the disk the method at [7] is adopted. Hence the image neighbourhood centered at C is convolved with a Laplacian of exponentially increasing sizes and the response normalized by the scale factor (a exponential growth of 1.2 is utilized). The radius of the Laplacian at the scale where the maximum response is ontained is the estimate of the disks radius.

The input to the method is a grayscale version of the original image. The output is represented as the center C and the radius ρ, of the optical disk. This information is stored in an output file and available to the user interface, upon image loading. Indicative results of the overall process are shown in Fig. 10.

Figure 10: Optic disk detection. Original image with the result of optic disk detection, superimposed by a yellow circle. The radius of the circle, visualizes the estimated disk size.

5 Integrated application

The functionalities provided by the software are integrated within a Graphical User Interface (GUI). Beside the conventional logistic operations (i.e. select file or interest, export measurement results to file, etc), the main goals of this GUI are:

- to support interactive measurements of medical professionals on regions and vessels of interest in an ergonomic way and to
- facilitate recovery from shortcomings of the automatic image processing algorithms and allow the expert user to modify their outcome.

The first goal is served by using an online representation of the detected vessels. In this way, the user can indicate vessels and vessel segments of interest with ease.

At the same time, the GUI provides a visualization of this representation indicating to the medical professional, the detected vessels on the original image. Using this visualization as feedback, the GUI supports the second goal by providing an editor where inaccurate vessel detection results can be corrected.

5.1 Measurement of a vessel segment

A typical operation during the inspection of the image by the medical professional is the measurement of the mean vessel width along a segment of a vessel.

The user provides two input points, with the mouse, to indicate the beginning and end of the segment. The two points closest to the user input, let $SP1$, $SP2$, of the skeleton are then found. The system then retrieves the skeleton points of the segment in between points $SP1$ and $SP2$ along with the estimated widths at these points and averages them.

The following algorithm is employed to find the points of a segment between two user provided points. It takes as input the coordinates of the two user-selected points, the skeleton image and the associated estimations of vessel width at the skeleton points. Its output is a skeleton segment and the mean vessel along its extent.

- For the user selected points (*UP1*, *UP2*) get the closest skeleton points (*SP1*, *SP2*), using an exhaustive search on all skeleton points based on distance.
- Starting from point *SP1*, consider all paths through the skeleton points that lead to *SP2*. Do this by starting a recursion for each neighbour of *SP1* and collecting each visited point in a separate list, for each recursion; typically there will be only two neighbours, but there can be more if *SP1* happens to be a branch point. End the recursion in either of the following cases:
  - an endpoint has been reached, or
  - point *SP1* is encountered, or
  - point *SP2* is encountered.

  Disregard results from the two first cases, as they do not lead to *SP2* they correspond to "dead-ends" and circles, respectively. In the first case, detect endpoints as skeleton points with only one neighbor. Save the result into memory in case *SP2* is encountered.
- In case no path reaching point *SP2* is found, the selected points are not connected and notify the user. Typically, there will be only one result (point list) connecting *SP1* and *SP2*. If more that one recursions encounter *SP2*, there exist multiple paths from *SP1* to *SP2*. In this case, select the shortest one, based on the number of points in the corresponding list.

Indicative results of the overall process are shown in Fig. 11.

5.2 Measurement of multiple vessel segments

A common measurement required by medical professionals is the measurement of average width of multiple non-branching vessels, at a range of distances from the center of the optical disk. The detection of these vessels is visualized in Fig. 12, where the first and last distances of the range are displayed as two concentric circles. Vessels which do not branch are detected and their estimated boundaries superimposed. Depending on user selection a number of statistics on the width of these vessels can be estimated (see Sec. 5.3).

The following algorithm takes as input the (i) center and radius of the optical disk, (ii) the range of distances to perform the measurement within (iii) the skeleton image S. Its output is the skeleton points within this range of distances organized per vessel along with the associated width measurements for each skeleton point. According to medical requirements, each return vessel segment initiates from the circle with the smaller radius and terminates until it reaches the larger radius or if a junction point is reached. Typically, the distance range is set so that the smaller and larger circles

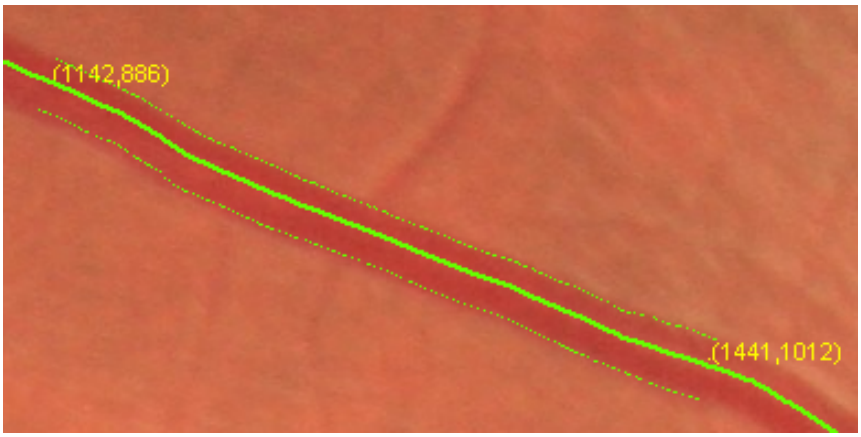

Figure 11: Width measurement of a vessel segment. The figure shows a detail of an original image, with the skeletonization result superimpose. The user indicates two points, their coordinates displayed in the above image. In the figure, the estimated boundaries for the indicated vessel segment are displayed (superimposed as green dots), providing a visualization of the information that the width measurement is based upon.

have radii that are 2 and 3 times longer than the radius of the optic disk.

- Collect all the skeleton points, *SKPDisk1*, of on the circumference of the smaller circle, based on their distance from the image center
- For each point *startSKP* of *SKPDisk1*:
  - Retrieve the path (*SKP_path*) of connected skeleton points starting from *startSKP*:
    - get the adjacent skeleton points, *adjacentSKpnts*
    - from *adjacentSKpnts*, select the point, *nextSKP*, with the larger distance from the center
    - initialize list *SKP_path* with the 2 skeleton points (*startSKP*, *nextSKP*)
    - expand recursively *SKP_path* following non-visited adjacent pixels and collecting visited points into *SKP_path*, until:
      - a junction point is reached, or
      - a skeleton point with distance from center $\geq$ *Disk2Radial* is reached, or
      - an endpoint has been reached and the path cannot be further expanded any more.
  - Discard the collected path in case:
    - of reaching a point which belongs as well to another path (vessels intersect in the image), or
    - of containing a point with distance from center $\leq$ *Disk1Radial* (case of a vessel intersecting the smaller circle in 2 different points)
  - Retrieve the corresponding width estimates for *SKP_path* and create a new vessel entry with both into *TotalPathsToExternalDisk*
- Sort *TotalPathsToExternalDisk* clockwise, by calculating the polar coordinates of the initial point of each vessel relative to center, in order to provide an intuitive enumeration of vessels in the display.

5.3 Statistical results

The calculation of a statistical measure M, introduced in [10, 11], has been integrated into the system and is available through the GUI. In brief this measure calculates the ratio of two quantities, $CRAE$ and $CRVE$, which are determined by measurement on the arteries and veins detected in the region of interest, respectively. The mean widths of vein and artery segments within the region of interest are collected in two separate lists, namely $ArteriolePathsWidths$ and $VenulePathsWidths$. Then

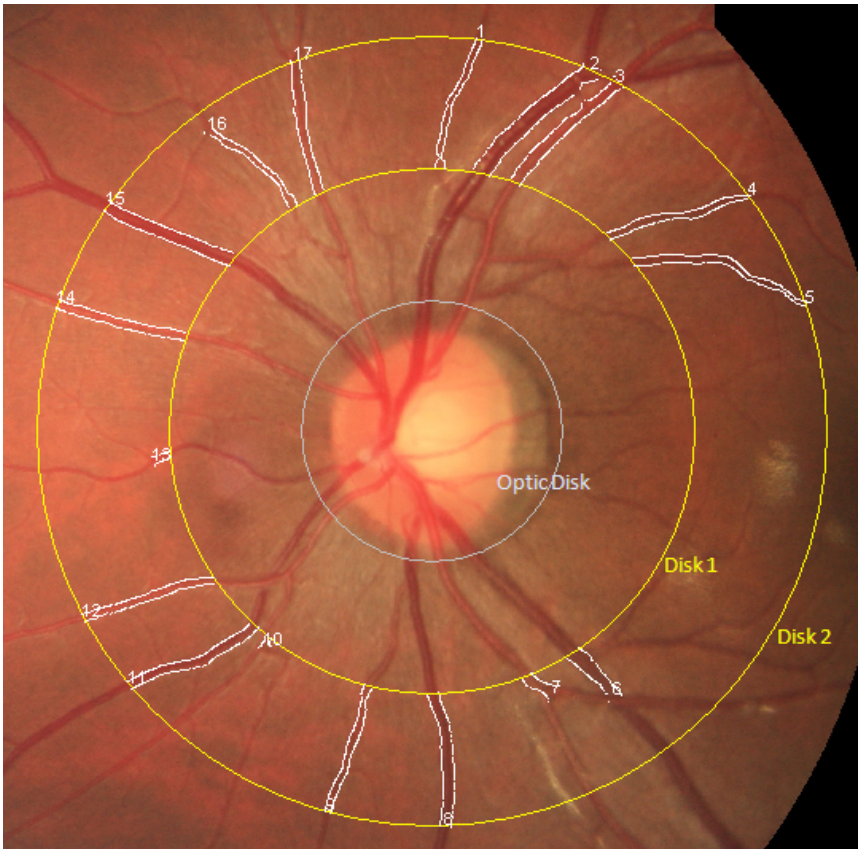

Figure 12: Detection and clockwise enumeration of vessel segments within the region between two concentric circles. Vessels 6,7 and 15 are terminated at their first branching point.

$CRAE$ (Central Retinal Arterial Equivalent) is computed as:

$$CRAE = \sqrt{.87 \cdot W_a^2 + 1.01 \cdot W_b^2 - .22 \cdot W_a \cdot W_b - 10.76}, \quad (1)$$

where W_b is the median value of $ArteriolePathsWidths$ and W_a is the value in $ArteriolePathsWidths$ exactly before the median. Correspondingly, $CRVE$ (Central Retinal Venous Equivalent) is computed as:

$$CRVE = \sqrt{.72 \cdot W_a^2 + .91 \cdot W_b^2 + 450.05}, \quad (2)$$

where now W_b is the median of $VenulePathsWidths$ and W_a the value in $VenulePathsWidths$ exactly before W_b.

Such measures require the characterization of vessels, as to if they are veins or arteries. Hence, a GUI component (see Fig. 13 is provided to facilitate this characterization by the medical professional. This component allows the user to select the type of each vessel (Arteriole / Venule) through a multiple choice radio button. The color of the vessels (red for arterioles, blue for venules) is automatically updated in the main window. A choice is also provided to "ignore" a particular vessel from statistical measurements. In addition the user is able to correct the mean width estimates for a vessel, in case of erroneous segmentation, by pressing a corresponding "fix" check-box and typing the fixed width in the corresponding text-input field.

5.4 Semi-automatic preprocessing

Using the user interface the operator of the software can edit the output of vessel segmentation, by adding and deleting pixels in the binary segmentation image. Correspondingly the skeleton representation and width measurements are updated with the new segmentation as input. The purpose of implementing this functionality is threefold:

1. Improve measurements by correcting for shortcomings of segmentation. That is, to provide the user a way to correct for remaining errors in the segmentation image and update measurement results accordingly.

2. Perform targeted measurements. Cases of interest include those where interest is focused on a particular vessel and its branches are required to be ignored. The medical professional can in this case delete segments that are to be ignored for measurements.

3. Acquire ground truth results. Using the interface a medical professional can have a basis to correct the result of segmentation and provide ground truth, regarding segmentation in retinal images.

Increasing the automation in processes 1 and 2 above is the topic of current and future work. The edited segmentations and corresponding results are stored in separate files.

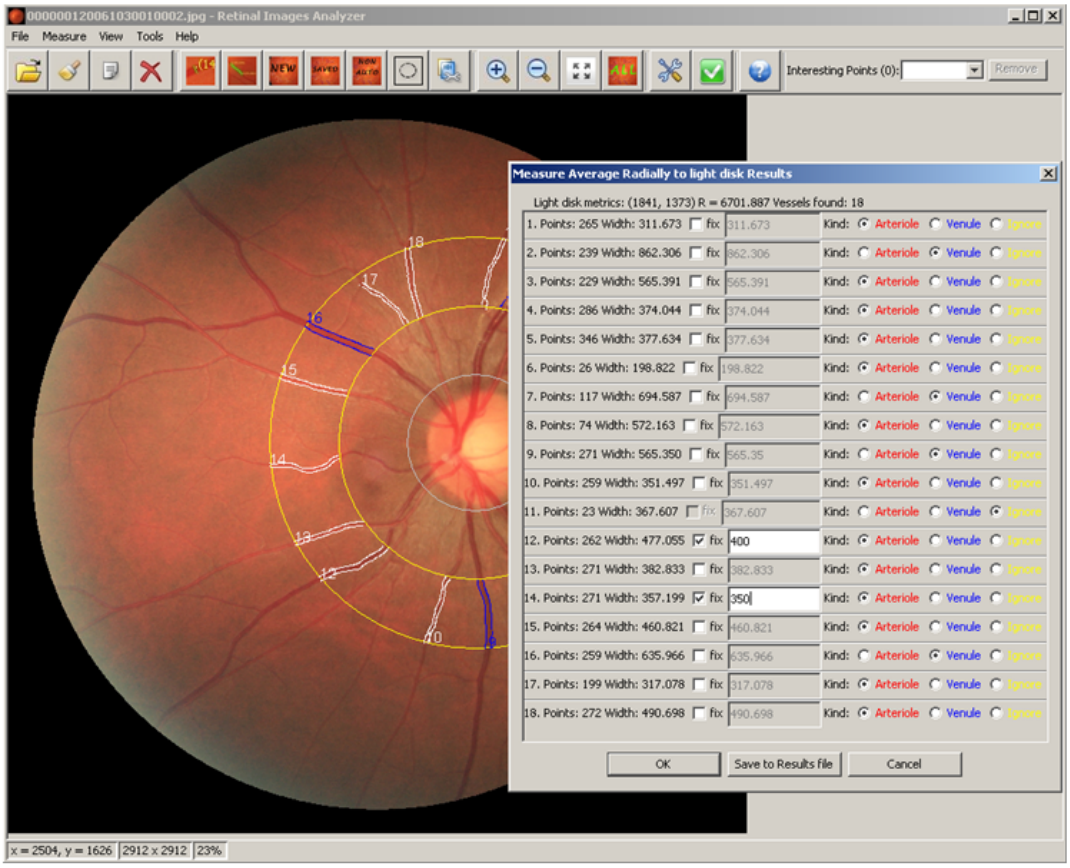

Figure 13: Characterization of vessel segments. The figure shows the vessel segments detected within a region of interest and a secondary characterization panel. The enumerated vessel segments in the image, are represented as a list in the panel. Using this panel the user can characterize each segment as to if it is a venule (indicated with blue outline in the image) or an arteriole (indicated by white outline). The user can also select which segments will be considered in the measurement. The analytical quantities of the measurement and the cumulative result can be exported to a file.

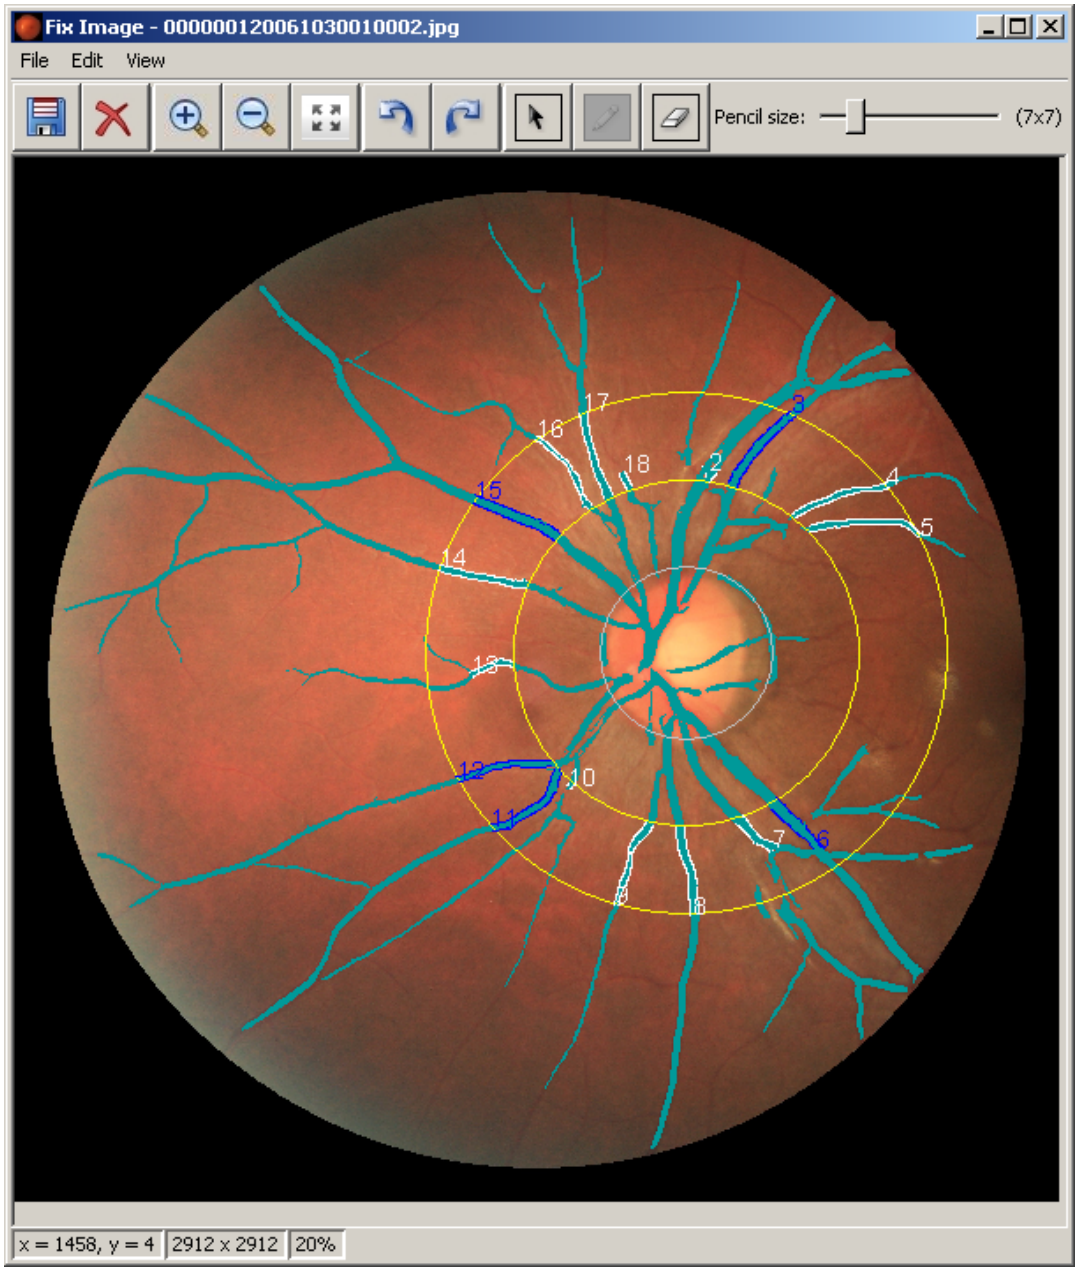

Figure 14: Segmentation result editing. The segmentation result is superimposed to the original image. The user can add and delete pixels with the, size-configurable, pencil and eraser tools respectively. The toolbar on the top provides conventional operations that facilitate the process, i.e. zoom-in, store intermediate result, multiple-level undo, etc.

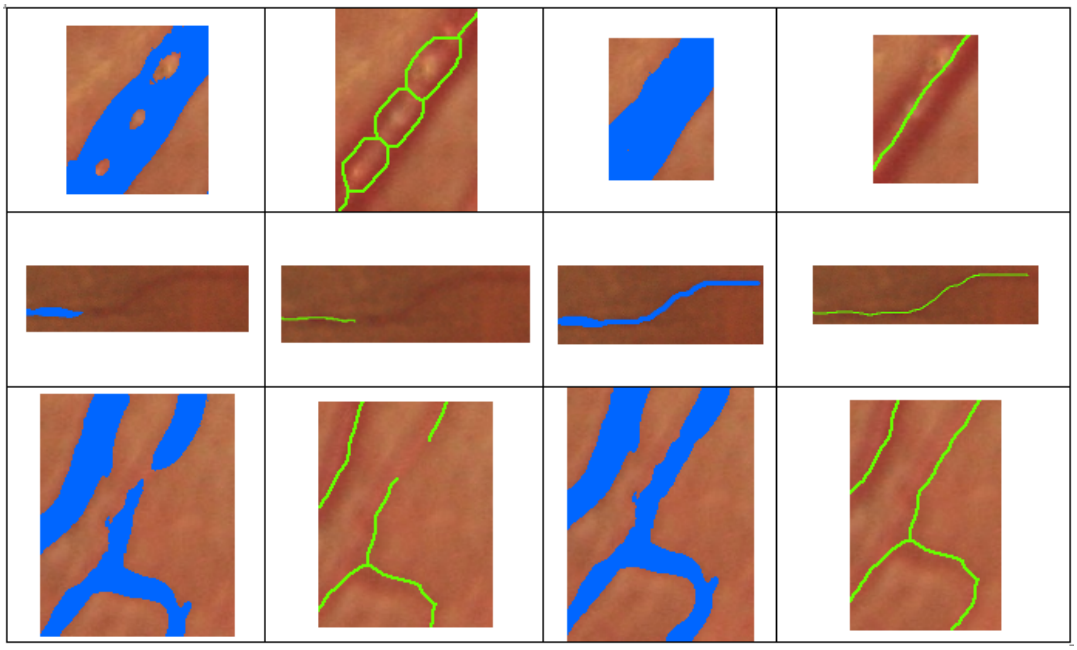

Figure 15: Segmentation and corresponding skeleton, before and after editing. Columns from left to right show, original segmentation, corresponding skeleton, edited segmentation, and its skeletonization.

An additional user interface panel that is optionally invoked facilitates the process (see Fig. 14). The examples below illustrate cases of use.

Example 1. Adding pixels. Typically used to correct falsely negatively classified pixels, as non-vessel. In Fig. 15, three characteristic examples are shown. In the first (top row), the segmentation exhibits holes in the segmented part within vessel structure vessel. As a consequence the result of the skeletonization is contains spurious vessel branches and segments. In the second (middle row), a detected vessel terminated earlier than the real one. By editing the vessel is expanded to a more accurate size. In the third (bottom row), a segmentation gap spuriously splits a vessel in two parts and, as a consequence, the skeleton represents the same vessel in different segments. The user can group the segments by filling the segmentation.

Example 2. Deleting pixels. Isolating a vessel of interest from branches. In Fig. 16, one characteristic example is shown. To focus measurement upon a vessel of interest, pixels corresponding to the branches of the vessel can be deleted, in a segmentation that isolates the vessel, or vessel arrangement of interest.

Example 3. Vessel width editing. In the example of Fig. 17, the segmentation represents the vessel with a wider extent than the veridical. By thinning the vessel, a more accurate segmentation and width measurements are achieved.

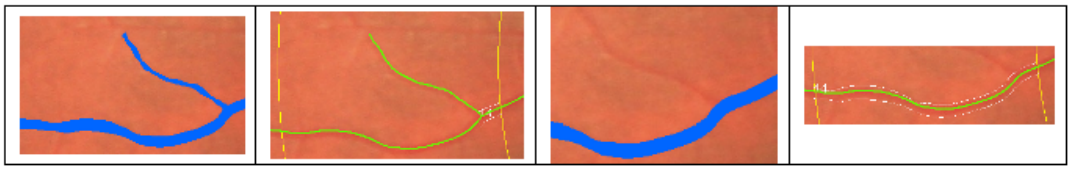

Figure 16: Segmentation and corresponding skeleton, before and after editing. Shown from left to right, original segmentation, corresponding skeleton, edited segmentation, and its skeletonization.

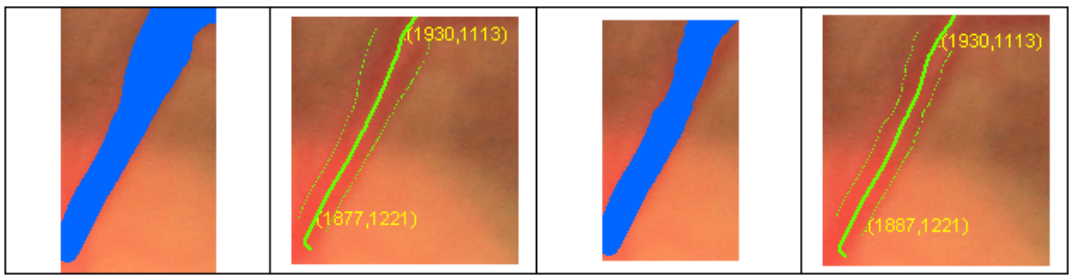

Figure 17: Segmentation and corresponding width measurement, before and after editing. The two images on the left show the original segmentation (leftmost), and the corresponding skeleton and the estimated vessel boundaries (middle-left). The two images on the right, show the edited segmentation (middle-right), and the updated skeleton boundary estimations (rightmost).

5.5 Application scenario and output

A characteristic scenario of use is presented in this subsection, indicating its context of use.

Initially the user can browse multiple files and select one or more image files to load their metrics information to the system. Such files have been precomputed either as a direct result of the employed image processing methods, or be an edited version of these results stored by the user. Having precomputed the measurement results accelerates user interaction and decreases user waiting. Large file collection are handled in batch mode by a pertinent software utility. A text file for each image, with coordinates of skeleton points, vessel width and coordinates of intersection points with vessel boundaries, is saved to disk.

Once the image is loaded, the user can select specific points on the vessels and measure their width online, using the method in Sec. 4.3. Another online measurement operation is the measurement of the mean vessel width along a vessel segment as described in Sec. 5.1. In both cases, the collected measurements can be saved for reviewing or exported in a text file for future processing.

The estimated location and size of the optic disk play a central role in the computation of region based statistics (see Sec. 4.4). Thus, the estimated center and radius are visualized on the display, superimposed on the original image as a circle. In case of inaccurate estimation, the user is able to adjust the estimated center by dragging this circle on the display and / or enter a new value for its radius. The computation of these statistics, then, is fully automatic and is invoked by a button on the GUI. The results are displayed on screen, but also exported in a text file.

Several other utility functionalities are performed through the user interface, such as the saving of data file with measurements, as well as, the reviewing and editing of older measurement files. In the display the user, can view the image and the calculated measurements and skeleton points using magnification options as well as rulers that indicate metric spatial scales (see Fig. 18). When in magnification a GUI component indicates the currently visible image segment, with respect to the total image (see Fig. 18).

All of the above functionalities are available through a toolbar of GUI buttons (see Fig. 19. Finally, the user interface provides auxiliary configuration panel where the user can modify the calibration so as to provide the correspondence of pixels to μm, as well as, other preferences such as font size for displaying the results.

6 Summary and future work

In this report, an application that increases the automation of vessel measurement in fundoscopy images is presented. This application employs an image segmentation algorithm, along with preprocessing and skeletonization techniques in order to extract a representation of vessels. In addition, techniques that analyze this representation and measure vessel width are introduced and adapted appropriately, in order to provide

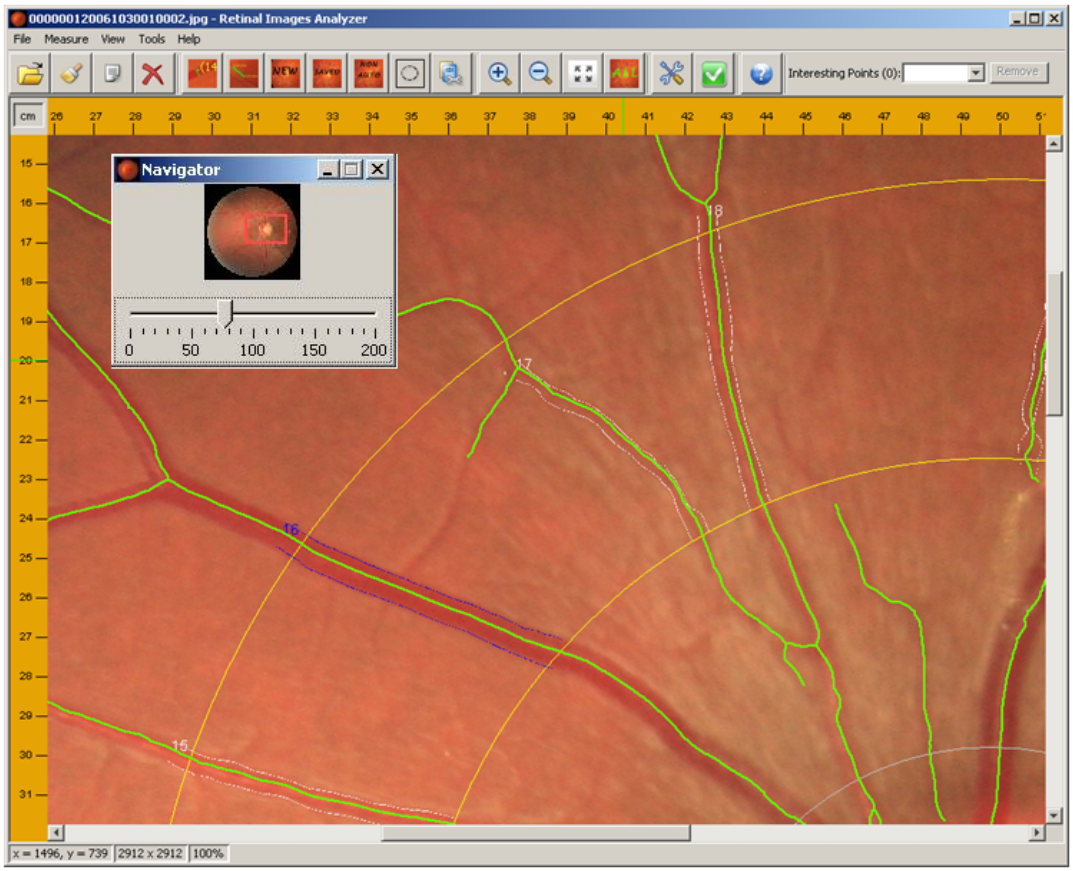

Figure 18: Image navigation utility panel of the user interface. Using this panel the user can display detailed views of the vessel representation and inspect them in context to the original image.

Figure 19: User interface toolbar. A toolbar in the graphical user interface, provide direct access to the main functionalities of the system.

measurements according to particular measurement protocols. The above functionalities are integrated through a Graphical User Interface, which assists the medical professional to perform measurements in an ergonomic fashion and apply targeted measurements according to medical protocols. This interface provides also functionalities that allow the medical professional to edit the vessel segmentation result and update the corresponding measurements, in order to recover from segmentation errors.

Future work will be pursued along two research avenues. The first regards the improvement of the segmentation algorithm, in order to obtain more accurate measurement results. Besides algorithmic improvements we also aim to address the fact that fundoscopy images a curved surface (the retina) and, therefore, vessel widths that are measured in the corresponding images have to be appropriately interpreted. The second regards the registration of fundoscopy images of the same patient that were acquired across large time intervals (i.e. 6 to 12 months). The goal is to enhance the functionalities of this application in a way that provides medical professionals with the capability of automatically comparing vessel measurements, thus offering a valuable tool in the monitoring, diagnosis, and estimation of the condition of the cardiovascular system.

Acknowledgements

This work has been supported by the FORTH-ICS internal RTD Programme 'Ambient Intelligence and Smart Environments'. We thank the Hypertension Unit, of the 2^{nd} Propedeutic Department of Internal Medicine, of the Hippokration Hospital, of the Aristotle University of Thessaloniki for posing the retinal vessel measurement problem as an application of interest, providing pertinent data, as well as, their insightful guidance into the requirements of medical professionals and contemporary diagnostic protocols. From this unit, we thank Areti Triantafyllou and Stella Douma for their contribution in the above collaboration. We also thank Vaggelis Sakkalis and Kostas Marias for continuing and fruitfull discussions on retinal image segmentation and preprocessing topics.

References

[1] Grosso A., Veglio F., Porta M., Grignolo F., and Wong T. Hypertensive retinopathy revisited: some answers, more questions. *British Journal of Ophthalmology*, 89(12):1646–1654, December 2005.

[2] Y. Cheng. Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):790–799, 1995.

[3] A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale vessel enhancement filtering. *Medical Image Computing and Computer Assisted Interventions, MICCAI98 (Boston, USA)*, 1496:130–137, 1998.

[4] Stromland K., Hellstrom A., and Gustavsson T. Morphometry of the optic nerve and retinal vessels in children by computer-assisted image analysis of fundus photographs. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 233(3):150–153, 1995.

[5] Hubbard L., Brothers R., King W., Clegg L., Klein R., Cooper L., Sharrett A., Davis M., and Cai J. Methods for evaluation of retinal microvascular abnormalities associated with hypertension/sclerosis in the atherosclerosis risk in communities study. *Ophthalmology*, 106(12):2269–2280, 1999.

[6] Lam L., S. Lee, and C. Suen. Thinning methodologies-a comprehensive survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):879, September 1992.

[7] T. Lindeberg. Detecting salient blob-like image structures and their scales with a scale-space primal sketch: A method for focus-of-attention. *International Journal of Computer Vision*, 11:283–318, 1993.

[8] A. Mendonca and A. Campilho. Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction. *IEEE Transactions on Medical Imaging*, 25(9):1200 – 1213, September 2006.

[9] Chapman N., Witt N., Gao X., Bharath A., Stanton A., Thom S., and A. Hughes. Computer algorithms for the automated measurement of retinal arteriolar diameters. *British Journal of Ophthalmology*, 85(1):74–79, 2001.

[10] J. Parr and G. Spears. Mathematic relationships between the width of a retinal artery and the widths of its branches. *Americal Journal of Ophtalmology*, 77:472–477, 1974.

[11] J.C. Parr and G.F. Spears. Mathematic relationships between the width of a retinal artery and the widths of its branches. *Am J Ophthalmol*, 77(4):478–83, 1974.

[12] Newsom R., Sullivan P., Rassam S., Jagoe R., and Kohner E. Retinal vessel measurement: comparison between observer and computer driven methods. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 230(3):221–225, 1992.

[13] Rassam S., Patel V., Brinchmann-Hansen O., Engvold O., and Kohner E. Accurate vessel width measurement from fundus photographs: a new concept. *British Journal of Ophthalmology*, 78(1):24–29, 1994.

[14] Wong T., Hubbard L., Klein R., Marino E., Kronmal R., Sharrett A., Siscovick D., Burke G., and Tielsch J. Retinal microvascular abnormalities and blood pressure in older people: the cardiovascular health study. *British Journal of Ophthalmology*, 86(9):1007–1013, 2002.

[15] X. Zabulis, A. A. Argyros, and D. P. Tsakiris. Lumen detection for capsule endoscopy. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3921–3926, 2008.

[16] K. Zuiderveld. Contrast limited adaptive histogram equalization. *Graphics gems IV*, pages 474–485, 1994.