CHAPTER 4

BACKGROUND

4.1 PREPROCESSING OPERATIONS

Retinal image preprocessing consists of detection of poor image quality, correction of non-uniform luminosity, color normalization and contrast enhancement.

4.1.1 Detection of Poor Image Quality

There are two kinds of quality problems in the fundus images used in this research: noise pixels and pixels whose color is distorted. Both seem to exist in regions where illumination has been inadequate. Since illumination is usually adequate in the center of the image, poor image quality regions are located near the edge of the fundus. Regions with poor image quality may cause errors in abnormality detection. There are three ways to solve image quality problems. The first solution is to repair incorrect regions. The second solution is not to process bad areas, in other words, anatomical landmarks are not searched from poor quality regions. However, in the second solution there is a risk that a fundus image is wrongly segmented. If the exclusive way is used, it is important that the algorithms provide information on the fundus image parts which are not processed. The third solution is to reject the whole fundus image if the area of poor image quality regions is too large. In the third solution the rejected fundus images should be examined manually. In this research, mainly the first and the second solutions are used. There are always some noise pixels in every fundus image due to inadequate illumination in some parts of the fundus. If those noise pixels are not in clusters but separated from each other, they can be easily removed by using a small median filter. However, if the lack of illumination has been significant in some compact regions, there may exist too many noisy pixels or pixels whose color is distorted, and the regions cannot be easily repaired. It was decided during the research that if there exist regions of low image quality in a fundus image, the image is not rejected until the proportion of bad quality regions is relatively large. The decision of segmentation of anatomical features is made according to regions of good image quality and the user of the algorithms is informed of regions having inadequate image quality.

4.1.2 Illumination Correction

One important issue in fundus images is that retina is not a plane surface and therefore light doesn't have a uniform distribution, producing images with non-uniform illumination and different contrast areas. The goal of illumination correction is to remove uneven illumination of the image caused by sensor defaults (vignetting), non uniform illumination of the scene, or orientation of the surface. In photography and optics, vignetting is a reduction of an image's brightness or saturation at the periphery compared to the image center. There are several causes of vignetting. Sidney F. Ray (2002) distinguishes the five types:

- Mechanical vignetting
- Optical vignetting
- Natural vignetting

A fourth cause is unique to digital imaging:

- Pixel vignetting

A fifth cause is unique to analog imaging:

- Photographic film vignetting

Improper scene illumination as well as non ideal acquisition conditions due to, for example, misadjusted imaging system can introduce severe distortions into the resulting image. These distortions are usually perceived as smooth intensity variations across the image. With such unevenness, subsequent image processing like image registration, segmentation, or pattern recognition may be substantially complicated; therefore, the correction of illumination inhomogeneities is highly desirable.

Illumination correction based on background subtraction is one of the widely used methods. There are two major types of background subtraction techniques depending on whether the illumination model of the images is given as an additional image or not. These are prospective correction and retrospective correction. Illumination variation in a fundus image can be eliminated by subtracting the background estimation from the original image or by dividing the original image by an estimated background. In this work, the method of luminosity correction used is based on a segmentation of background pixels. Here, we use a method of luminosity correction that is based on Joshi and Sivaswamy (2008) by the segmentation of background pixels and subsequent computation of luminosity function based only on the background image. The advantage of this approach is that it does not produce any ringing effect.

4.1.3 Colour Channel Separation

A colour image is composed of three colour channels, *viz* red, green and blue and hence known as an RGB image. A colour fundus image and its component channels are shown below. Figure 4.1 (a) shows the original image, Figure 4.1 (b) shows the red channel, Figure 4.1 (c) shows the green channel and Figure 4.1 (d) shows the blue channel.

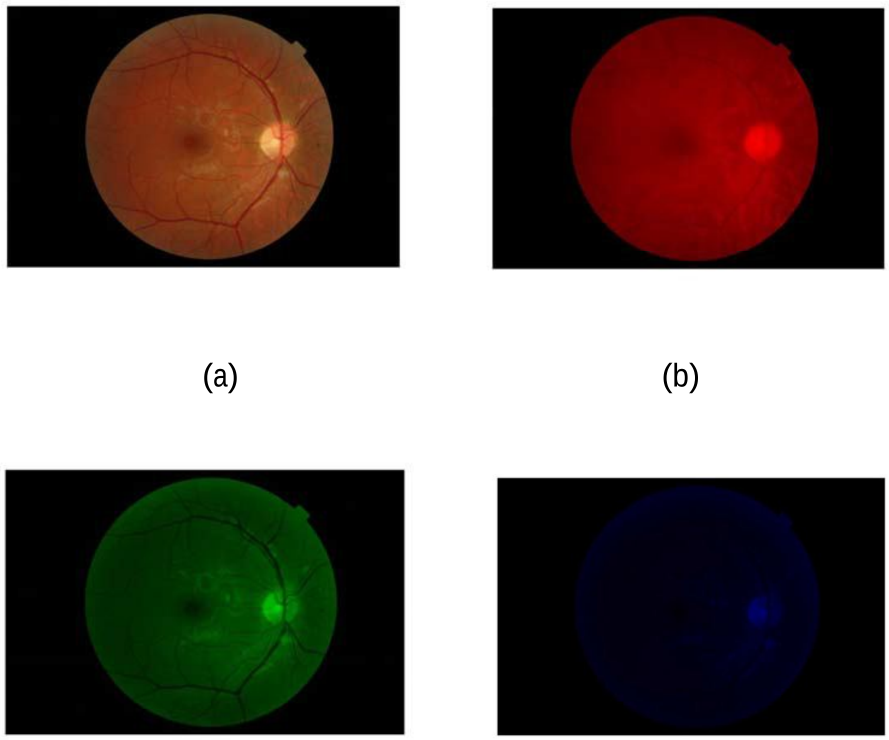

(a) (b)

(c) (d)

Figure 4.1 RGB image and its components

4.1.4 Tophat transform

The efficiency of the tophat transform was introduced by Meyer in 1979 for Cytology applications. In mathematical morphology and digital image processing, tophat transform is an operation that extracts small elements and details from given images. There exist two types of top-hat transform, the white top-hat transform and the black top-hat transform. The white top-hat transform is defined as the difference between the input image and its opening by some structuring element. The black top-hat transform is defined dually as the difference between the closing and the input image. The

Top-hat operation is used to separate foreground from background. Mathematically, it can be defined as Equation 4.1 and 4.2, where " $\bullet$ " represents morphological closing, " $\circ$ " represents morphological opening, X is the original image and B represents the structuring element .

$$White\ Tophat(X) = X - X \circ B \tag{4.1}$$

$$Bottomhat\ orBlack\ Tophat(X) = X \bullet B - X \tag{4.2}$$

The white top-hat transform returns an image, containing those objects or elements of an input image that are smaller than the structuring element (i.e., places where the structuring element does not fit in), and are brighter than their surroundings (Serra 1982, 1988).

The black top-hat returns an image, containing the objects or elements that are smaller than the structuring element, and are darker than their surroundings. It is sometimes called the bottomhat transform. Bottom hat transform subtracts the original image from a morphologically closed version of the image. It can be used to find intensity troughs in an image.

The size, or width, of the elements that are extracted by the top-hat transforms can be controlled by the choice of the structuring element B. The bigger the latter, the larger the elements extracted. Both top-hat transforms are images that contain only non-negative values at all pixels. The peaks extraction in grey level images is done by the white tophat transform while the valleys extraction is done by black tophat transform. Top-hat transforms are used for various image processing tasks, such as feature extraction, background equalization, image enhancement, and others.

4.1.5 Contrast Enhancement using Tophat Transforms

The tophat transform or the white tophat transform, as seen from section 4.1.4, is defined as the difference between the original image and its opening . The opening of an image is the collection of foreground parts of an image that fit a particular structuring element. The bottomhat transform (black tophat transform) is defined as the difference between the closing of the original image and the original image. The closing of an image is the collection of background parts of the image that fits the chosen structuring element. The tophat image contains the peaks of objects that fit the structuring element. In contrast, the bottomhat image shows the gaps between the objects of interest. Thus a combination of these transforms can be used for contrast enhancement. To maximize the contrast between the objects and the gaps that separate them from each other, the bottomhat image can be subtracted from the 'original image +tophat image'. That is for an image X and B structuring element, the contrast enhanced image Y is,

$$WTH = White\ Tophat(X) = X - X \circ B$$

$$BTH = Bottomhat(X) = X \bullet B - X$$

$$Y = (X + WTH) - BTH \tag{4.3}$$

4.1.6 Histogram Equalization

Histogram of an image is a plot of the number of occurrences of grey levels in the image against the grey level values. The histogram provides a convenient summary of the intensities in an image, but it is unable to convey any information regarding spatial relationships between pixels (Jayaraman et al. 2009). The histogram provides more insight about image contrast and brightness. The histogram of a dark image will be clustered towards the lower grey level while it will be clustered towards the higher grey level in the case of a bright image. For a low contrast image, the histogram will not be spread equally, that is, the histogram will be narrow but for a high contrast image, there will be an equal spread in the grey level. Image brightness and contrast appearance may be improved by modifying the histogram.

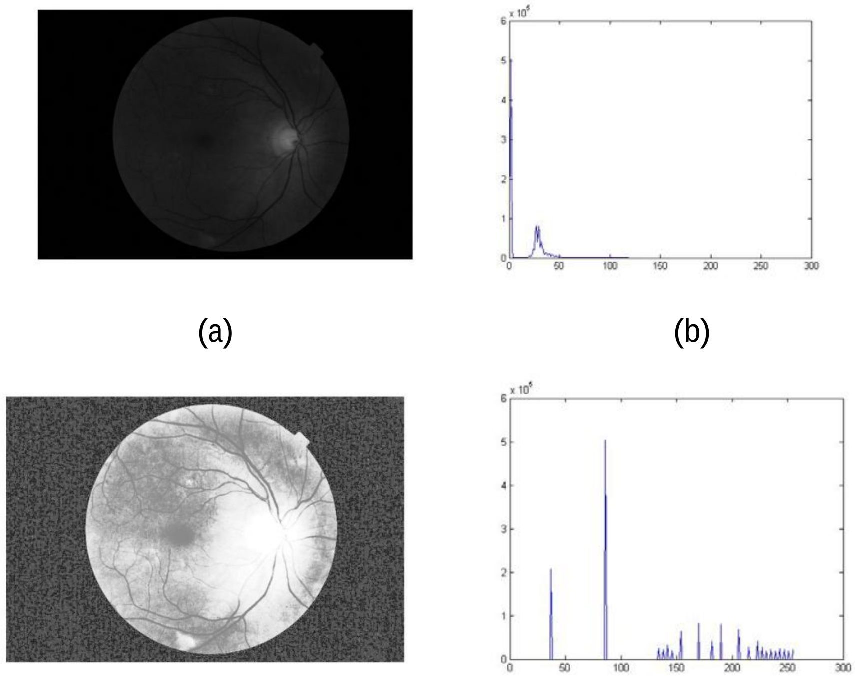

(a) (b)

(c) (d)

Figure 4.2 Histogram equalization

Histogram equalization is used to equally distribute the number of pixels between grey levels. It stretches or compresses the image. Histogram equalization reassigns the brightness values of pixels based on the image histogram. Histogram equalization enhances the contrast of images by transforming the values in an intensity image, or the values in the colour map of an indexed image, so that the histogram of the output image approximately matches a specified histogram. Histogram equalization results in a histogram of the equalized image as flat as possible and produces more visually pleasing results. Histogram matching allows us to specify which pixel values we want to express or depress, in other words it lets us to choose any histogram shape. Figure 4.2 shows the effect of histogram equalization for a grey level image. Figure 4.2(a) shows the original image and Figure 4.2(b) shows its corresponding histogram, while Figure 4.2 (c) shows the histogram equalized image and Figure 4.2 (d) shows its corresponding histogram.

4.1.7 Contrast Limited Adaptive Histogram Equalization (CLAHE)

Adaptive histogram equalization (AHE) is an image processing technique used to improve contrast in images. It differs from ordinary histogram equalization in the respect that the adaptive method computes several histograms, each corresponding to a distinct section of the image, and uses them to redistribute the lightness values of the image. It is therefore suitable for improving the local contrast of an image and bringing out more detail.

However, AHE has a tendency to over amplify noise in relatively homogeneous regions of an image. A variant of adaptive histogram equalization called contrast limited adaptive histogram equalization (CLAHE) prevents this by limiting the amplification (http://en.wikipedia.org/wiki/ adaptive_histogram_equalization). Ordinary histogram equalization uses the same transformation derived from the image histogram to transform all pixels. This works well when the distribution of pixel values is similar throughout the image. However, when the image contains regions that are significantly lighter or darker than most of the image, the contrast in those regions will not be sufficiently enhanced.

In its simplest form, each pixel is transformed based on the histogram of a square surrounding the pixel. The derivation of the transformation functions from the histograms is exactly the same as for ordinary histogram equalization: The transformation function is proportional to the cumulative distribution function (CDF) of pixel values in the neighborhood.

CLAHE differs from ordinary adaptive histogram equalization in its contrast limiting. This feature can also be applied to global histogram equalization, giving rise to contrast limited adaptive histogram equalization (CLAHE), which is rarely used in practice. In the case of CLAHE, the contrast limiting procedure has to be applied for each neighborhood from which a transformation function is derived. CLAHE was developed to prevent the over amplification of noise that adaptive histogram equalization can give rise to. This is achieved by limiting the contrast enhancement of AHE. The contrast amplification in the vicinity of a given pixel value is given by the slope of the transformation function. This is proportional to the slope of the neighborhood cumulative distribution function (CDF) and therefore to the value of the histogram at that pixel value. CLAHE limits the amplification by clipping the histogram at a predefined value before computing the CDF. This limits the slope of the CDF and therefore of the transformation function. The value at which the histogram is clipped, the so-called clip limit, depends on the normalization of the histogram and thereby on the size of the neighborhood region. Common values limit the resulting amplification to between 3 and 4.

It is advantageous not to discard the part of the histogram that exceeds the clip limit but to redistribute it equally among all histogram bins. The redistribution will push some bins over the clip limit again, resulting in an effective clip limit that is larger than the prescribed limit and the exact value of which depends on the image. If this is undesirable, the redistribution procedure can be repeated recursively until the excess is negligible.

Adaptive histogram equalization in its straightforward form presented above, both with and without contrast limiting, requires the computation of a different neighborhood histogram and transformation function for each pixel in the image. This makes the method very expensive computationally.

Interpolation allows a significant improvement in efficiency without compromising the quality of the result. The image is partitioned into equally sized rectangular tiles (64 tiles in 8 columns and 8 rows is a common choice.). A histogram, CDF and transformation function is then computed for each of the tiles. The transformation functions are appropriate for the tile center pixels. All other pixels are transformed with up to four transformation functions of the tiles with center pixels closest to them, and are assigned interpolated values. Pixels in the bulk of the image are bilinearly interpolated, pixels close to the boundary are linearly interpolated, and pixels near corners are transformed with the transformation function of the corner tile. The interpolation coefficients reflect the location of pixels between the closest tile center pixels, so that the result is continuous as the pixel approaches a tile center. This procedure reduces the number of transformation functions to be computed dramatically and only imposes the small additional cost of linear interpolation.

4.2 BITPLANE DECOMPOSITION

The grey level of each pixel in a digital image is stored as one or more bytes in a digital computer. When the grey level is represented in a single byte, it is called an 8 bit image, representing grey level values in the range 0 to 255. The bit-plane representation of an 8 bit image is as shown in Figure 4.3.

Slicing a digital image into its bit planes is useful for analyzing the relative importance played by each bit of the image. Instead of highlighting gray level images, highlighting the contribution made to total image appearance by specific bits is examined in this . In an 8 bit gray level image, the image is represented by 8 bits. The image is composed of 8, 1-bit planes ranging from bit plane 0 (LSB)to bit plane 7 (MSB). It is to be noted that the change in Most Significant Bit (MSB) significantly changes the value encoded by the byte while a change in Least Significant Bit (LSB) does not change the encoded grey level much. In terms of 8-bits, plane 0 contains all lowest order bits in the bytes comprising the pixels in the image and plane 7 contains all higher order bits. Thus bitplane decomposition of an 8 bit image yields eight binary images.

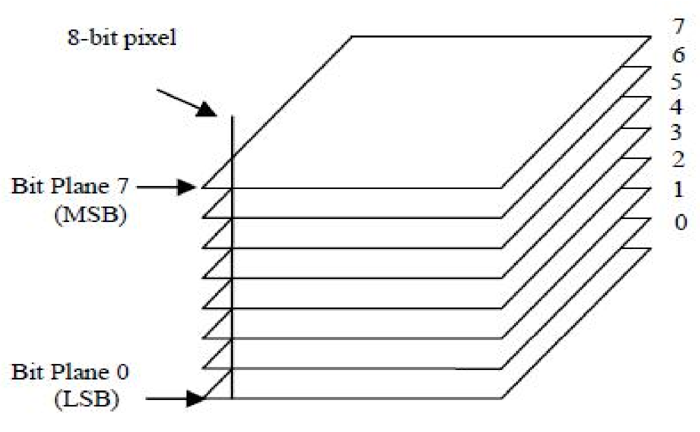

Figure 4.3 Bitplane decomposition

The following are the important applications of bitplane decomposition (Jayaraman et al. 2009).

- Convert a grey level image to a binary image

- Represent an image with fewer bits and compress it to a smaller size
- Enhance the image by focusing.

4.3 MATHEMATICAL MORPHOLOGY

Mathematical Morphology (MM) is considered as the science of appearance, shape and organization (Jayaraman et al. 2009). MM is a theory and technique for the analysis and processing of geometrical structures, based on set theory, lattice theory, topology, and random functions. MM is most commonly applied to digital images, but it can be employed as well on graphs, surface meshes, solids, and many other spatial structures. Topological and geometrical continuous-space concepts such as size, shape, convexity, connectivity, and geodesic distance, can be characterized by MM on both continuous and discrete spaces (http://en.wikipedia.org/wiki/Mathematical_ morphology). MM is also the foundation of morphological image processing, which consists of a set of operators that transform images according to the above characterizations. MM was originally developed for binary images, and was later extended to grayscale functions and images.

In digital image processing, mathematical morphology deals with non-linear processes which can be applied to an image to remove details smaller than a certain reference shape, usually called the structuring element. The basic idea in binary morphology is to probe an image with a simple, pre-defined shape, drawing conclusions on how this shape fits or misses the shapes in the image. This simple "probe" is called structuring element, and is itself a binary image (i.e., a subset of the space or grid).

The most common morphological operations used in image processing are dilation, erosion, opening and closing. The basic operations are shift-invariant (translation invariant) operators strongly related to *Minkowski* addition. The images obtained after bit plane decomposition are binary images, which are thus ideal for performing morphological operations.

Dilation is an operation in which the binary image is expanded from its original shape. The amount of expansion is controlled by the structuring element. The dilation process is similar to convolution, in which the structuring element is reflected and shifted from left to right and then from top to bottom. In this process, any overlapping pixels under the centre position of the structuring element are assigned with 1 or black values.

If X is the reference image and B is the structuring element, the dilation of X by B is represented as

$$X \oplus B = \{Z|[(\hat{B})z \cap X] \subseteq X\} \tag{4.4}$$

where $\hat{B}$ is the image B rotated about the origin. When an image X is dilated by a structuring element B, the outcome element Z would be that there will be at least one element in B that intersects with an element in X.

Erosion is a thinning operation that shrinks an image. The extent by which shrinking takes place is determined by the structuring element. Here, if there is a complete overlapping with the structuring element, the pixel is set white or 0. The erosion of X by B is given as

$$X \ominus B = \{Z|[(B)z] \subseteq X\} \tag{4.5}$$

In erosion, the outcome element Z is considered only when the structuring element is a subset or equal to the binary image X.

Opening operation is performed by first doing an erosion, followed by a dilation. Opening smoothens the inside of object contours, breaks narrow strips and eliminate thin portions of the image. It is mathematically represented as

$$X \circ B = (X \ominus B) \oplus B \tag{4.6}$$

Closing operation does the opposite of opening. It is dilation followed by erosion. Closing fills small gaps and holes in a single pixel object. The closing process is represented by

$$X \bullet B = (X \oplus B) \ominus B \tag{4.7}$$

Closing operation protects coarse structures, closes small gaps and rounds off concave corners. Morphological operations are widely used in the detection of boundaries in a binary image.

For an image X , the following can be applied to obtain a boundary image.

$$Y = X - (X \ominus B)$$

$$Y = (X \oplus B) - X$$

or

$$Y = (X \oplus B) - (X \ominus B) \tag{4.8}$$

where, the operator '$\oplus$' denotes dilation , '$\ominus$' denotes erosion and '$-$ ' indicates set theoretical subtraction. Most binary morphological operations have natural extensions to gray scale processing. Some, like morphological reconstruction have applications that are unique to gray scale images , such as peak filtering.

4.4 RETINAL IMAGE DATABASES

Several publicly available databases for different purposes have been provided by universities and research organizations all over the world as described below. We have used the MESSIDOR database , DRIVE database and the STARE database for the evaluation of our algorithms.

4.4.1 The MESSIDOR Database

The name MESSIDOR is evolved from "Methods to evaluate segmentation and indexing techniques in the field of retinal ophthalmology". The MESSIDOR database has been established to facilitate studies on computer-assisted diagnoses of diabetic retinopathy. The data included in this database can be used, free of charge, for research and educational purposes. The 1200 eye fundus color numerical images of the posterior pole for the MESSIDOR database were acquired by 3 ophthalmologic departments using a color video 3CCD camera on a Topcon TRC NW6 non-mydriatic retinograph with a 45 degree field of view. The images were captured using 8 bits per color plane at 1440 x 960, 2240 x 1488 or 2304 x 1536 pixels. A total of 800 images were acquired with pupil dilation (one drop of Tropicamide at 0.5%) and 400 without dilation. The 1200 images are packaged in 3 sets, one per ophthalmologic department. Each set is divided into 4 zipped sub sets containing each 100 images in TIFF format and an Excel file with medical diagnoses for each image.

For each image, two diagnoses, retinopathy grade and risk of macular edema, have been provided with the dataset. These diagnoses were obtained by medical experts following a grading scheme. These diagnoses were considered as the reference standard for the performance analysis. According to the reference standard, a total of 546 images were classified as normal and 654 as presenting signs of DR, specifically 153 with retinopathy grade 1, 247 with retinopathy grade 2 and 254 with retinopathy grade 3. Additionally, 974 images do not show risk of macular edema; whereas 75 and 151 images presented risk 1 and 2 of macular edema, respectively. Information about patients were removed in order to ensure patient privacy.

All of the images contained in the database were used for making actual clinical diagnoses. To ensure the protection of patient privacy, information that might allow the identity of a patient to be reconstructed has been concealed. The database is available for direct download for educational and research purposes from http://messidor.crihan.fr.

4.4.2 The DRIVE Database

The DRIVE (Digital Retinal Images for Vessel Extraction) database, consists of a total of 40 color fundus photographs (Staal et al. 2004). The DRIVE database has been established to enable comparative studies on segmentation of blood vessels in retinal images. All images have been de-identified, they were stripped from all individually identifiable information and processed in such a way that this information cannot be reconstructed from the images. The photographs were obtained from a diabetic retinopathy screening program in the Netherlands. The screening population consisted of 453 subjects between 31 to 86 years of age. Each image has been JPEG compressed, which is common practice in screening programs. Of the 40 images in the database, 7 contain pathology, namely exudates, hemorrhages and pigment epithelium changes.

The images were acquired using a Canon CR5 non-mydriatic 3CCD camera with a 45 degree field of view (FOV). Each image is captured using 8 bits per color plane at 768 X 584 pixels. The FOV of each image is circular with a diameter of approximately 540 pixels. The set of 40 images was divided into a test and training set both containing 20 images. Three observers manually segmented a number of images. All observers were trained by an experienced ophthalmologist. The first observer segmented 14 images of the training set while the second observer segmented the other 6 images. The test set was segmented twice resulting in a set X and Y. Set X was segmented by both the first and second observer (13 and 7 images respectively) while set Y was completely segmented by the third observer.

It is available for the public and research community from their website: http://www.isi.uu.nl/Research/Databases/DRIVE/.

4.4.3 The STARE Database

The name STARE is an abbreviation of structured analysis of the retina. The STARE Project was conceived and initiated in 1975 by Michael Goldbaum at the University of California, San Diego, and has been funded continuously by the National Institutes of Health (U.S.A.) since 1986. During this time, over thirty people have contributed to the project, with backgrounds ranging from medicine to science and engineering. Images and clinical data have been provided by the Shiley Eye Centerat the University of California, San Diego, and by the Veterans Administration Medical Center in San Diego. It is available for educational and research purposes from their website http://www.ces.clemson.edu/~ahoover/stare/.

4.4.4 ARIA Online

The ARIA online retinal image dataset is a joint research project between St Paul's Eye Unit, Royal Liverpool University Hospital Trust, Liverpool, UK and Ophthalmology, Clinical Sciences, University of Liverpool, Liverpool UK. The ARIA project aims to provide an automated image capture and image analysis platform capable of predicting individuals at risk of eye disease and that can be used at the point of image capture in the community.

This directory contains images collected by members of staff of St Paul's Eye Unit and the University of Liverpool as part of the ARIA project. All subjects were adults. All fundus images were taken using a Zeiss FF450+ fundus camera and originally stored as uncompressed TIFF files. The images were converted to compressed JPG files for on-line publication on their web site. All photographs were taken at a 50 degree field width, and all images are in colour. Additionally, trained image analysis experts have traced out the blood vessels in the images and these are also given in sub-directories. The optic disk and fovea, where relevant, have also been outlined in separate sets of image files. The data is organised into three categories, namely, that from age-related macular degeneration (AMD) subjects, healthy control-group subjects, and diabetic subjects.

4.4.5 The DIARETDB Database

This consists of DIARETDB0 and DIARETDB1. This is a public database for benchmarking diabetic retinopathy detection from digital images. The main objective of the design has been to unambiguously define a database and a testing protocol which can be used to benchmark diabetic retinopathy detection methods. By using this database and the defined testing protocol, the results between different methods can be compared. The databases can be freely downloaded and used for scientific research purposes. DIARETDB0 and DIARETDB1 are copyrighted in 2006 and 2007 respectively by Tomi Kauppi and his team.

The DIARETDB0 database consists of 130 color fundus images of which 20 are normal and 110 contain signs of the diabetic retinopathy (hard exudates, soft exudates, micronaneuyrysms, hemorrhages and neovascularization). Images were captured with a 50 degree field-of-view digital fundus camera with unknown camera settings. The database correspond to practical situations, and can be used to evaluate the general performance of diagnosis methods. This data set is referred to as "calibration level 0 fundus images".

The DIARETDB1 database consists of 89 colour fundus images of which 84 contain at least mild non-proliferative signs (Microaneurysms) of the diabetic retinopathy, and 5 are considered as normal which do not contain any signs of the diabetic retinopathy according to all experts who participated in the evaluation. Images were captured using the same 50 degree field-of-view digital fundus camera with varying imaging settings. The data correspond to a good (not necessarily typical) practical situation, where the images are comparable, and can be used to evaluate the general performance of diagnostic methods. This data set is referred to as "calibration level 1 fundus images". The DIARETDB1 database can be accessed in the web from http://www2.it.lut.fi/project/imageret/diaretdb1/.

4.4.6 The HEI-MED Dataset

The Hamilton Eye Institute Macular Edema Dataset (HEI-MED) (formerly DMED) is a collection of 169 fundus images in JPEG format to train and test image processing algorithms for the detection of exudates and diabetic macular edema (Giancardo et al. 2012). The images have been collected as part of a telemedicine network for the diagnosis of diabetic retinopathy currently developed by the Hamilton Eye Institute, the Image Science and Machine Vision Group at ORNL with the collaboration of the Université de Bourgogne.

Each image of the dataset was manually segmented by Dr. Edward Chaum (an expert ophthalmologist from HEI). He identified all the exudates and other bright lesions such as cotton wool spots, drusens or clearly visible fluid leakage occurring in the fundus. There was no distinction between hard and soft exudates because this differentiation is prone to errors and does not provide a clear clinical advantage for the diagnosis. In addition to the images and the ground truth, the dataset provide other anonymous clinical metadata about the patients, the optic nerve's manually identified location, the machine segmented vasculature (employing the method of Zana and Klein) and a Matlab class to seamlessly access all the data and metadata without having to deal with the internal format of the files. It is available for download from their website http://vibot.u-bourgogne.fr/luca/heimed.php.

4.4.7 The ORIGA-(light) Online Depository

ORIGA stands for Online Retinal Fundus Image database for Glaucoma Analysis and research. It is described by Zhang et al. in their paper as an online depository, which aims to share clinical ground truth (GT) retinal images with the public and to provide open access for researchers to benchmark their computer-aided segmentation algorithms. An in-house image segmentation and grading tool is developed to facilitate the construction of ORIGA(-light). Currently, ORIGA(-light) contains 650 retinal images annotated by trained professionals from Singapore Eye Research Institute. A wide collection of image signs, critical for glaucoma diagnosis, are annotated. ORIGA(-light) is available for online access upon request.

4.4.8 Retinopathy Online Challenge (ROC)

The ROC aims to help patients with diabetes through improving computer aided detection and diagnosis (CAD) of diabetic retinopathy. ROC facilitates the translation of diabetic retinopathy CAD into clinical practice by:

1. enabling any medical image analysis research group to develop diabetic retinopathy CAD algorithms by offering a training set of retinal images with reference standard provided by internationally accepted retinal experts.

2. evaluating the output of a diabetic retinopathy CAD algorithm in a uniform manner on a supplied test set, allowing algorithms to be compared both to other algorithms and retinal experts.

3. organizing meetings and workshops at international conferences to compare CAD systems, following the paradigm of revolution through competition.

Currently, they have released a first data set, aimed at CAD of microaneurysms and dot hemorrhages. On this site, interested research groups and companies can register a team, download the data and submit results. The ROC microaneurysm dataset is available online and this competition website remains active and open for new submissions (Niemeijer et al. 2010).

In addition to this competition, powerful online annotation tools available on this site enable experts to browse image data on-line, annotate it, and in future test their reading skills and compare their detection performance with the reference standard and with that of CAD. ROC can be accessed from http://roc.healthcare.uiowa.edu/.

4.4.9 The VARIA Database

VARPA is a research group affiliated to the Department of Computer Science of the Faculty of Informatics of the University of Coruña. One of the most active fields of work in the VARPA Group is the ophtalmology, in particular the analysis of eye fundus images (retinal images). These images are mostly granted by the Complejo Hospitalario Universitario of Santiago (CHUS). VARIA database is maintained by this group.

The VARIA database is a set of retinal images used for authentication purposes (Ortega et al. 2009). The database currently includes 233 images, from 139 different individuals. The images have been acquired with a TopCon non-mydriatic camera NW-100 model and are optic disc centered with a resolution of 768x584. The database distribution includes a directory with the images and a index.txt file indicating which images are from each user.

The VICAVR database is a set of retinal images used for the computation of the A/V Ratio. This database currently includes 58 images. The images have been acquired with a TopCon non-mydriatic camera NW-100 model and are optic disc centered with a resolution of 768x584. The database includes the caliber of the vessels measured at different radii from the optic disc as well as the vessel type (artery/vein) labeled by three experts.