LAPPEENRANTA UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF INFORMATION TECHNOLOGY

RETINAL IMAGE ANALYSIS USING MACHINE VISION

The topic of the master's thesis has been accepted in the Departmental council of the Department of Information Technology, on September 15, 2004.
Supervisors: Professor Heikki Kälviäinen, Dr. Tech. Joni-Kristian Kämäräinen

Lappeenranta, June 6, 2005

Markku Kuivalainen
Huopatehtaankatu 2 E 4
53900 Lappeenranta
Tel. +358 407 395 850
Markku.Kuivalainen@lut.fi

ABSTRACT

Lappeenranta University of Technology

Department of Information Technology

Markku Kuivalainen

Retinal Image Analysis Using Machine Vision

Master's thesis

2005

96 pages, 53 figures, 2 tables, and 1 appendix.

Supervisors: Professor Heikki Kälviäinen and Dr. Tech. Joni-Kristian Kämäräinen

Keywords: machine vision, image processing, pattern recognition, diabetic retinopathy, fundus image, retina, retinal image

The topic of this thesis is studying how lesions in retina caused by diabetic retinopathy can be detected from color fundus images by using machine vision methods. Methods for equalizing uneven illumination in fundus images, detecting regions of poor image quality due to inadequate illumination, and recognizing abnormal lesions were developed during the work. The developed methods exploit mainly the color information and simple shape features to detect lesions. In addition, a graphical tool for collecting lesion data was developed. The tool was used by an ophthalmologist who marked lesions in the images to help method development and evaluation. The tool is a general purpose one, and thus it is possible to reuse the tool in similar projects.

The developed methods were tested with a separate test set of 128 color fundus images. From test results it was calculated how accurately methods classify abnormal funduses as abnormal (sensitivity) and healthy funduses as normal (specificity). The sensitivity values were 92% for hemorrhages, 73% for red small dots (microaneurysms and small hemorrhages), and 77% for exudates (hard and soft exudates). The specificity values were 75% for hemorrhages, 70% for red small dots, and 50% for exudates. Thus, the developed methods detected hemorrhages accurately and microaneurysms and exudates moderately.

TIIVISTELMÄ

Lappeenrannan teknillinen yliopisto

Tietotekniikan osasto

Markku Kuivalainen

Silmänpohjakuvien analysointi konenäöllä

Diplomityö

2005

96 sivua, 53 kuvaa, 2 taulukkoa ja 1 liite.

Tarkastajat: Professori Heikki Kälviäinen ja TkT Joni-Kristian Kämäräinen

Hakusanat: konenäkö, kuvankäsittely, hahmontunnistus, diabeettinen retinopatia, silmänpohjakuva, verkkokalvo, retina

Keywords: machine vision, image processing, pattern recognition, diabetic retinopathy, fundus image, retina, retinal image

Diplomityössä tutkittiin diabeettiseen retinopatiaan liittyvien verkkokalvomuutosten tunnistamista konenäkömenetelmillä silmänpohjavärikuvista. Työssä kehitettiin menetelmät silmänpohjakuvien epätasaisen valaistuksen korjaamiseen, valaistuksen heikkoudesta johtuvien häiriöiden havaitsemiseen ja itse verkkokalvomuutosten tunnistamiseen. Kehitetyt menetelmät hyödyntävät pääasiassa väri-informaatiota ja yksinkertaisia muotoparametreja verkkokalvomuutosten tunnistuksessa. Työn alkuvaiheessa kehitettiin lisäksi graafisella käyttöliittymällä varustettu datankeräystyökalu, jonka avulla asiantuntija merkitsi käytetyistä silmänpohjakuvista löytyneet verkkokalvomuutokset. Datankeräystyökalu suunniteltiin yleiskäyttöiseksi, joten sitä voidaan käyttää tarvittaessa myös muissa vastaavantyyppisissä projekteissa.

Kehitetyt menetelmät testattiin 128 värillisellä silmänpohjakuvalla. Testin tuloksista määriteltiin menetelmien kyvyt luokitella verkkokalvomuutoksia sisältävät silmänpohjat epänormaaleiksi (sensitiivisyys) ja vastaavasti terveet silmänpohjat normaaleiksi (spesifisyys). Menetelmien sensitiivisyydet olivat verenvuodoille 92%, pienille punaisille muutostyypeille (mikroaneurysmat ja pienet verenvuodot) 73% ja eksudaateille (kovat ja pehmeät eksudaatit) 77%. Vastaavasti menetelmien spesifisyydet olivat verenvuodoille 75%, pienille punaisille muutostyypeille 70% ja eksudaateille 50%. Kehitetyt menetelmät löysivät siis verenvuodot hyvin sekä mikroaneurysmat ja eksudaatit tyydyttävästi.

Preface

This master's thesis was done in the Department of Information Technology at Lappeenranta University of Technology during autumn 2004 and spring 2005.

I would like to thank my supervisors, Professor Heikki Kälviäinen and Dr. Tech. Joni-Kristian Kämäräinen, for making it possible to work with an interesting topic and for giving support with the research when needed. I would thank the University of Kuopio, especially Valentina Kalesnykiene, for providing fundus images and medical knowledge for this research. I also express my gratitude to Tekes and the YTI Research Center for funding the research.

Lappeenranta, June 6, 2005

Markku Kuivalainen

Contents

Symbols and abbreviations

B	Blue color channel in RGB image
B_r	Relative blue color channel
F_n	False negative
F_p	False positive
G	Green color channel in RGB image
G_r	Relative green color channel
H	Hue channel in HSI or HSV image
I	Intensity channel in HSI image
k	Constant slope used in linear threshold equation
R	Red color channel in RGB image
R_r	Relative red color channel
S	Saturation channel in HSI or HSV image
T	Threshold value or thresholding function
t	Threshold value
t_0	Constant threshold value
T_n	True negative
T_p	True positive
V	Value (intensity) channel in HSV image
w	Mask coefficient
θ	Hue angle in degrees in HSI color space
ε	Error value
$e(x, y)$	Estimation function
$f(x, y)$	Discrete digital image
$f_{green}(x, y)$	Discrete green color channel (G)
$f_{red}(x, y)$	Discrete red color channel (R)
$g(x, y)$	Processed image
$i(x, y)$	Source illumination
$p(x, y)$	Local property of a pixel
$r(x, y)$	Reflectance component
$r_{green}(x, y)$	Reflectance component for green color channel
x_c	X-coordinate of a center point
y_c	Y-coordinate of a center point

$\oplus$	Morphological dilation
$\ominus$	Morphological erosion
$\circ$	Morphological opening
$\bullet$	Morphological closing

AND	Logical And Operation
DFT	Discrete Fourier Transform
FCM	Fuzzy C-means Clustering
HSI	Hue, Saturation, and Intensity; a color system
HSV	Hue, Saturation, and Value; a color system
IDFT	Inverse Discrete Fourier Transform
RAM	Random Access Memory
RGB	Red, Green, and Blue; a color system
SOM	Self-Organizing Map
SVM	Support Vector Machine
XOR	Logical Exclusive-Or Operation

1 Introduction

This research is a part of a larger project investigating how changes in the retina caused by certain diseases can be automatically detected from fundus images. The project involves, in addition to Lappeenranta University of Technology, the University of Kuopio, Mikkeli Polytechnic, and a few eye clinics. The project is funded by Tekes [1], the main public funding organization for research and development in Finland.

The project is divided into separate research topics, and the research presented in this master's thesis concerns the detection of diabetic retinopathy lesions. Another research site, located in the YTI Research Center in Mikkeli Polytechnic, develops methods for detecting blood vessels, optic disk, and location of the fovea. The Department of Ophthalmology at the University of Kuopio provides fundus images and lesion content of the images.

When the methods are fully developed, they will be combined together into an automatic fundus image analysis system. The automatic system is meant to be used in hospital level screening and monitoring to provide alerts of possible findings for further medical inspection.

Diabetes is a metabolic disorder where the blood glucose level has increased [2], [3]. There are two types of diabetes: Type 1 is insulin dependent and Type 2 non-insulin dependent diabetes [2], [4]. Most of the diagnosed patients in Type 1 diabetes are children or young persons, but the disease may also appear in adulthood [5]. Type 2 diabetes usually appears in middle-aged or elderly people, but nowadays it is also found among younger groups [6].

Almost 200 000 Finnish persons, almost 4% of the population, have been diagnosed to have either Type 1 or Type 2 diabetes [7]. In addition, it is also believed that there are at least 50 000 undiagnosed Type 2 diabetics in Finland [2]. According to specialists, the number of Type 2 diabetics will increase by 70% in the next ten years if no precautional actions are performed [2]. Type 1 diabetes is more common in Finland than in any other country [5]. The existence of Type 2 diabetes is not so high, but it is still the most common in Finland of the Nordic countries [2]. Diabetes (Type 1 and 2) is the most expensive national disease in Finland [8]: it costs about 850 000 - 950 000 euros a year, which is 11-12% of the total public health service costs in Finland [9]. Thus, diabetes is a very common and expensive disease in the Finnish national health care.

Diabetes can cause a variety of complications, especially when untreated. One of the most serious complications is kidney failure, whereas problems with feet and vision are the most common ones [10], [4].

1.1 Objectives and limitations

The main objective of this research was to develop reliable and accurate image processing and pattern recognition methods for automatic fundus image analysis. The main focus is on accurate and reliable detection of abnormal lesions, belonging to diabetic retinopathy, from color fundus images. The different diabetic retinopathy lesion types that are to be detected are microaneurysms, hemorrhages, soft exudates, and hard exudates. Also neovascularization was first involved in the research, but it was excluded due to a lack of resources. However, for consistency, neovascularization is described briefly and it is also discussed whether the other implemented algorithms could be used for finding neovascularizations.

The secondary objective was to develop a general tool for collecting medical image data. The user of the tool, for example a domain expert, should be able to efficiently and quickly mark lesions in the images so that the collected data can later be used by image processing experts. In this research, the tool was used by an ophthalmologist who marked abnormal lesions into fundus images. The images and their corresponding lesion information, or ground truth, were used for developing and testing the methods.

The time available for this research was limited to seven months. The relatively short time available had to be prioritized in order to cover at least the most important features.

1.2 Structure of the thesis

The rest of this thesis consists of six sections. Section 2 familiarizes the reader with the concept of machine vision and describes the techniques of machine vision used in this research and in other relative studies in the literature. Section 3 gives an introduction to diabetic retinopathy and describes the different abnormality lesion types of the disease. The section explains how the screening for diabetic retinopathy is done at the moment. The section also familiarizes the reader with related work presented in the literature. Section 4 describes how the fundus images of this research were acquired and how they were en-

hanced before abnormality detection. The methods for detecting abnormal lesions are presented in Section 5. The experiments and their results are presented in Section 6, and finally, conclusions are drawn in Section 7.

2 Machine vision

Gonzalez and Woods [11] explain that machine vision, or computer vision, has the ultimate goal *"to use computers to emulate human vision, including learning and being able to make inferences and take actions based on visual inputs"*. In addition to the visual band of the electromagnetic spectrum perceived by humans, imaging machines cover almost the entire electromagnetic spectrum from gamma to radio waves. Thus, digital image processing can be used in a wide and varied field of applications [11].

Machine vision is widely used in industrial quality control systems since machine vision applications do not have the foibles of human beings, such as tiredness and illness, and since machine vision-based systems are faster, more accurate and inexpensive than human employees. Machines can also be used in environments that are too dangerous for human beings. Machine vision can be used in security and surveillance applications, for example, in face detection and recognition and intruder monitoring systems. In medicine machine vision can be used for example to detect abnormal lesions or in radiation therapy planning where radiation has to be targeted exactly to a tumor without causing damage to the surrounding tissue.

Figure 1 shows the principle of a typical machine vision-based quality control system. A camera is connected to a computer having a proper machine vision software. The camera captures an image of each of the items to be inspected and sends the image to the computer. The machine vision software analyzes the image and decides whether the item is valid or not according to the domain knowledge stored in the computer. A separate rejecting system can be used for separating rejected items from passed items.

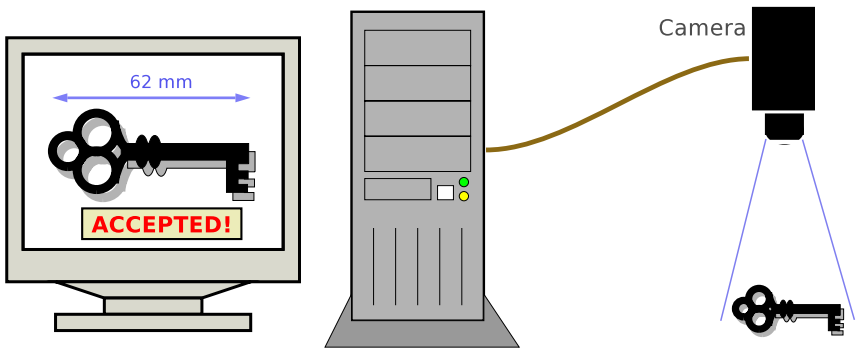

Figure 1: The principle of a typical machine vision application.

The phases of a machine vision application are shown in Figure 2. First, an image is cap-

tured by a some type of camera (still camera, video camera, spectral camera, etc.). Then the acquired image is pre-processed where the quality of the image is increased. After pre-processing the image is segmented so that objects of interest are separated from other objects and the background. Then the separated objects are classified into pre-defined classes. The results of the classification can be used when deciding what operations are performed for the objects in the image. Such an operation could be rejecting an object due to its low quality if the machine vision application is used in quality control.

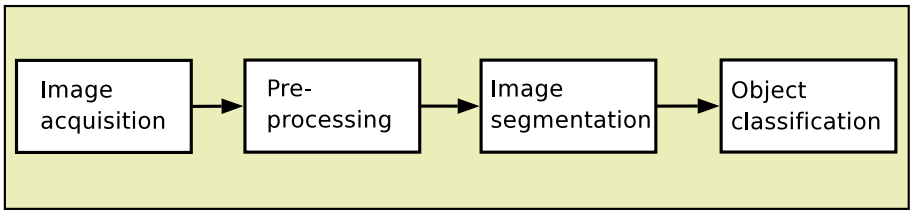

Figure 2: Phases of a machine vision application.

2.1 Image pre-processing

The aim of pre-processing is to increase the quality of an image by reducing the amount of noise appearing in the image and highlighting features that are used in image segmentation. Two typical techniques used in pre-processing are filtering and contrast enhancing. The following subsections describe both techniques.

2.1.1 Filtering

A digital gray-scale image may be defined as a two-dimensional function, $f(x, y)$, where x and y are discrete spatial coordinates and f is discrete gray-level of the image. Thus, a digital image consists of a finite number of elements having finite gray-levels. The discrete elements of a digital image are called picture elements or pixels.

Filtering is a common technique in digital image processing, used either for enhancing the quality of images or making them more suitable for further processing. Filtering can be used for example for sharpening blurry images or removing noise from images. Image filtering can be performed either in the spatial or frequency domain. Spatial filtering operates directly on the pixels that compose the image. In the frequency domain filtering the image is first converted to the frequency domain where the filtering is performed. The

conversion from the spatial to the frequency domain can be done for example by discrete Fourier transform (DFT). After filtering the image is converted back to the spatial domain for example by inverse discrete Fourier transform (IDFT).

Spatial filtering is based on a small mask that is simply moved from pixel to pixel in the image. At each pixel the result of the filter is calculated using a preset rule. Figure 3(a) shows a representation of a general 3 x 3 spatial filter mask where w_1, w_2, ..., w_9 are mask coefficients. In general, a filter mask can be of size m x n, where the width $m = 2a + 1$, height $n = 2b + 1$, and a and b are non-negative integers. For linear spatial filtering, for example smoothing and sharpening, the result image g at each image pixel (x, y) can be calculated as a sum of the products of the filter coefficients w and the corresponding image pixels f [11] as

$$g(x,y) = \sum_{s=-b}^{b} \sum_{t=-a}^{a} w(t,s) f(x+t, y+s). \tag{1}$$

Note that the pixels near the edge of the image (closer than the a pixels in the horizontal direction or closer than the b pixels in the vertical direction) cannot be filtered since the location of a pixel in the equation should have non-negative integers. There are several ways to handle the pixels near the image edge: they can be omitted, causing the result image to be smaller than the original one, or those pixels can be set for example to 0's.

w_1	w_2	w_3
w_4	w_5	w_6
w_7	w_8	w_9

(a)

$1/9$	$1/9$	$1/9$
$1/9$	$1/9$	$1/9$
$1/9$	$1/9$	$1/9$

(b)

Figure 3: 3 x 3 spatial filter masks: (a) General 3 x 3 spatial filter mask; (b) 3 x 3 mean filter mask.

Smoothing filters can be used for blurring images and reducing noise in the images. The output of smoothing is simply the average of pixels covered by the filter mask. Thus, smoothing filters are also called averaging filters or mean filters. Figure 3(b) represents a simple 3 x 3 mask used in smoothing.

In addition to the linear filters described above, order-statistics filters are also used in digital image processing. The term order-statistics means that the filters are non-linear spatial filters whose response is based on ordering the pixels covered by the filter mask. The best-known filter belonging to the order-statistics class is the median filter. Median filters are commonly used because they provide excellent noise reduction for certain ran-

dom noise types and they do not blur images as heavily as linear filters. The response for a median filter is the median value of the gray-level pixels covered by the filter mask. As an example of median filtering, suppose that pixels covered by a 3 x 3 filter mask have the gray-level values 30, 35, 5, 30, 40, 35, 70, 30, 35. The values are sorted as 5, 30, 30, 30, 35, 35, 35, 40, 70, where the median value is 35 since it is in the center of the sorted values.

Figure 4 shows an example of using mean and median filters. Figure 4(a) shows a gray-level image that is heavily corrupted by salt-and-pepper noise: 15% of total number of pixels were randomly replaced with white and black pixels. Figure 4(b) shows the result of 3 x 3 mean filtering where noise has been reduced but is still visible. Figure 4(c) represents the result of 3 x 3 median filtering where almost all noise has been removed, and thus, the quality of the image has appreciably increased.

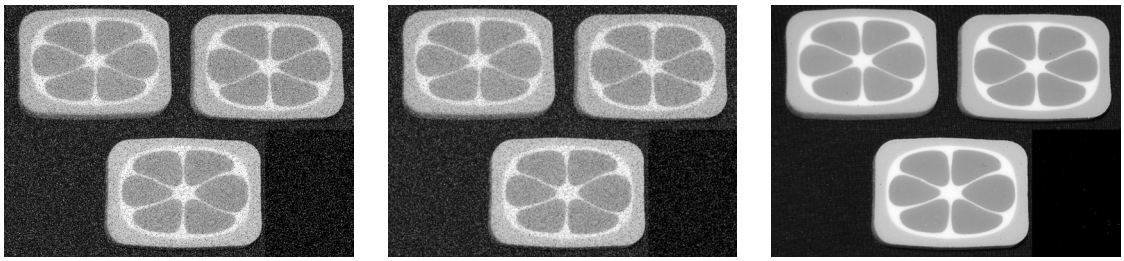

(a) (b) (c)

Figure 4: Example of spatial filtering: (a) Original image of size 1110 x 800 pixels corrupted by salt-and-pepper noise; (b) Result of noise reduction with a 3 x 3 mean filter; (c) Result of noise reduction with a 3 x 3 median filter.

2.1.2 Contrast enhancement

Images having low contrast can result from inadequate illumination, wrong lens aperture settings, or lack of dynamic range in the imaging sensor. Contrast enhancement increases the dynamic range of an image.

One technique for contrast enhancement is contrast-stretching transformation which maps the gray-levels of the original image to new levels as shown in Figure 5 [11]. The locations of points (r_1, s_1) and (r_2, s_2) define the shape of the transformation function. An example of contrast-stretching is shown in Figure 6. Figure 6(a) is an RGB image suffering of low contrast. Figure 6(b) is result of performing contrast-stretching transformation separately for each of the three color channels, red, blue, and green, in the original RGB image.

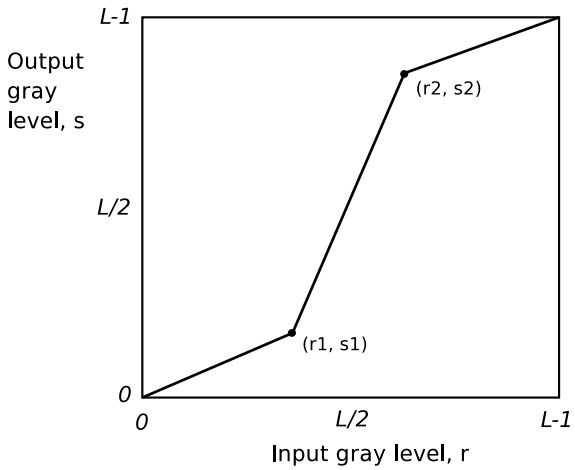

Figure 5: Mapping of gray-levels in contrast-stretching transformation. L is the number of gray-levels in the image.

(a) (b)

Figure 6: Example of contrast-stretching transformation: (a) Original RGB image; (b) Result of performing contrast stretching transformation to the original image.

Histogram equalization is another method for contrast enhancement. Figure 7(a) shows the histogram of the red color channel in the original RGB image in Figure 6(a). As seen in the histogram only a part of the gray-levels are used in the image. The histogram equalization method stretches the gray-levels so that the whole gray-scale is smoothly used. Figure 7(b) shows a result RGB image where the histogram equalization has been performed separately for each of the color channels. A histogram of the red color channel (R) in the result image is shown in Figure 7(c). As visible in the histogram, gray-levels are used more evenly after the histogram equalization.

Local contrast enhancement is also a histogram-based contrast enhancement method. In

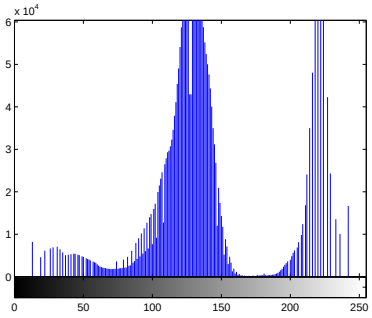

(a)

(b)

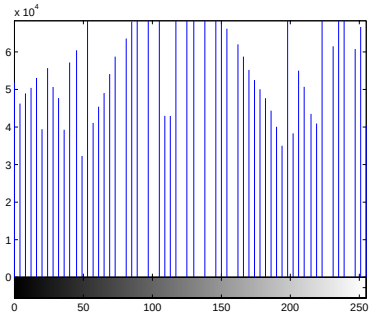

(c)

Figure 7: Example of histogram equalization: (a) Histogram of the red color channel (R) in Figure 6(a); (b) Result of histogram equalization; (c) Histogram of the red color channel in the result image.

local contrast enhancement an image is divided into sub-blocks and the histogram equalization is done separately for each of the blocks. After the histogram equalization the whole image is reproduced by combining the results of the sub-blocks with interpolation. The interpolation smoothly combines the sub-blocks into a whole image so that the junctions of the sub-blocks are not visible. Figure 8(a) shows the result of performing local contrast enhancement separately for each of the color channels in the image shown in Figure 6(a). A histogram of the red color channel in the result image is shown in Figure 8(b). As seen in the histogram, the medium gray-levels are used more than lowest or highest gray-levels in the result image. In this example image local contrast enhancement gave a superior result compared to contrast stretching transformation and histogram equalization, as the result of local contrast enhancement is the most natural and pleasant.

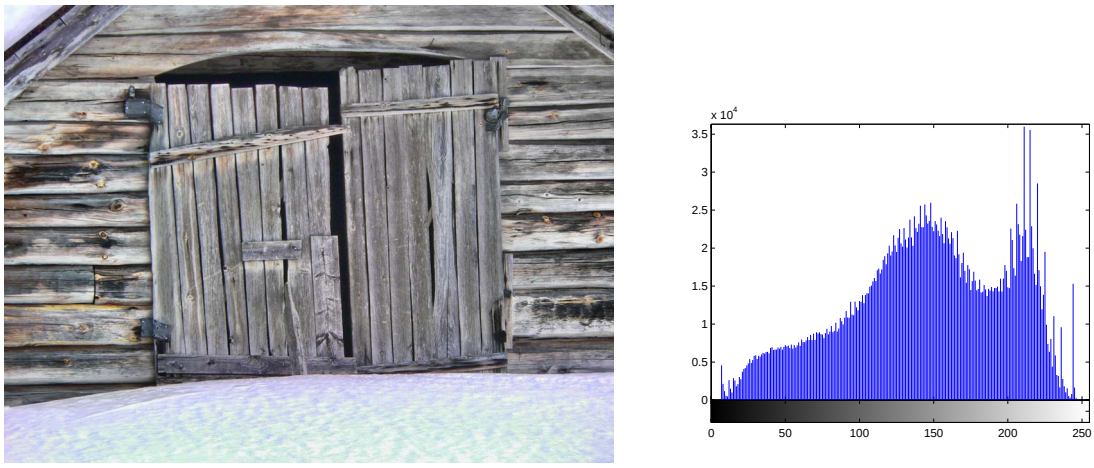

(a) (b)

Figure 8: Example of local contrast enhancement: (a) Result of performing local contrast enhancement for Figure 6(a); (b) Histogram of the result image.

2.2 Image segmentation

In traditional image processing both input and output are images. Techniques for enhancing images to be more suitable for visualizing purposes or further processing were described in Sub-section 2.1 above. In machine vision problems it is necessary to understand what objects are visible in the images being processed. Thus, information of objects of interest are extracted from the images. Segmentation is one of the most widely used steps in reducing images to information [12].

Segmentation divides an image into regions or objects. The level of division depends on the problem being solved, meaning that the segmentation of the image should stop when the objects of interest have been isolated. According to Gonzalez and Woods [11], the segmentation of nontrivial images is one of the most difficult tasks in image processing. They also remind that segmentation accuracy determines the eventual success or failure of computerized analysis procedures.

The algorithms for segmenting images are based on discontinuity or similarity in pixel values. In the discontinuity approach an image is partitioned on the basis of abrupt changes in intensity, such as edges, points, blobs, and lines in the image. In the similarity approach a set of pre-defined criteria is used for partitioning the image into regions that are similar according to the criterion set. Thresholding, region growing, and region splitting and merging are examples of similarity-based segmentation techniques.

In the next sub-sections thresholding, region growing, and blob detection are described with illustrative examples. The methods are first described for gray-scale images and then a separate sub-section explains color image segmentation. There are also other, more sophisticated image segmentation techniques described in the literature, but they are out of the scope of this thesis.

2.2.1 Thresholding

Thresholding is an important and widely used technique in image segmentation, because thresholding is effective and simple to implement. Thresholding is based on the assumption that the foreground (objects) can be separated from the background according to the gray-level values. In thresholding, pixels within a defined range are selected as belonging to the foreground whereas gray-levels outside the range are rejected to the background.

Thresholding can be considered as an operation involving tests against function T [11]:

$$T = T[x, y, p(x, y), f(x, y)] \tag{2}$$

where $f(x, y)$ is the gray-level of the pixel in location (x, y) and $p(x, y)$ is some local property of the pixel, for example, the mean gray-level of the pixel neighborhood. The thresholded image $g(x, y)$ can be defined as

$$g(x, y) = \begin{cases} 1 & \text{if } f(x, y) \geq T \\ 0 & \text{if } f(x, y) < T. \end{cases} \tag{3}$$

When T depends only on $f(x, y)$, the threshold is called global. If T also depends on $p(x, y)$, the threshold is called local. The threshold is called dynamic or adaptive if T depends also on spatial coordinates x and y.

Figure 9 shows an example of global thresholding. Figure 9(a) is a gray-scale image $f(x, y)$ where there are two light objects on a dark background. Figure 9(b) is a histogram of Figure 9(a). As seen in the histogram, the objects can be extracted from the background by selecting a threshold T that separates the two groups of histogram peaks. If a pixel at location (x, y) has value $f(x, y) \geq T$, the pixel belongs to an object; otherwise the pixel is a background pixel. Figure 9(c) shows the result of thresholding the image with threshold value 100.

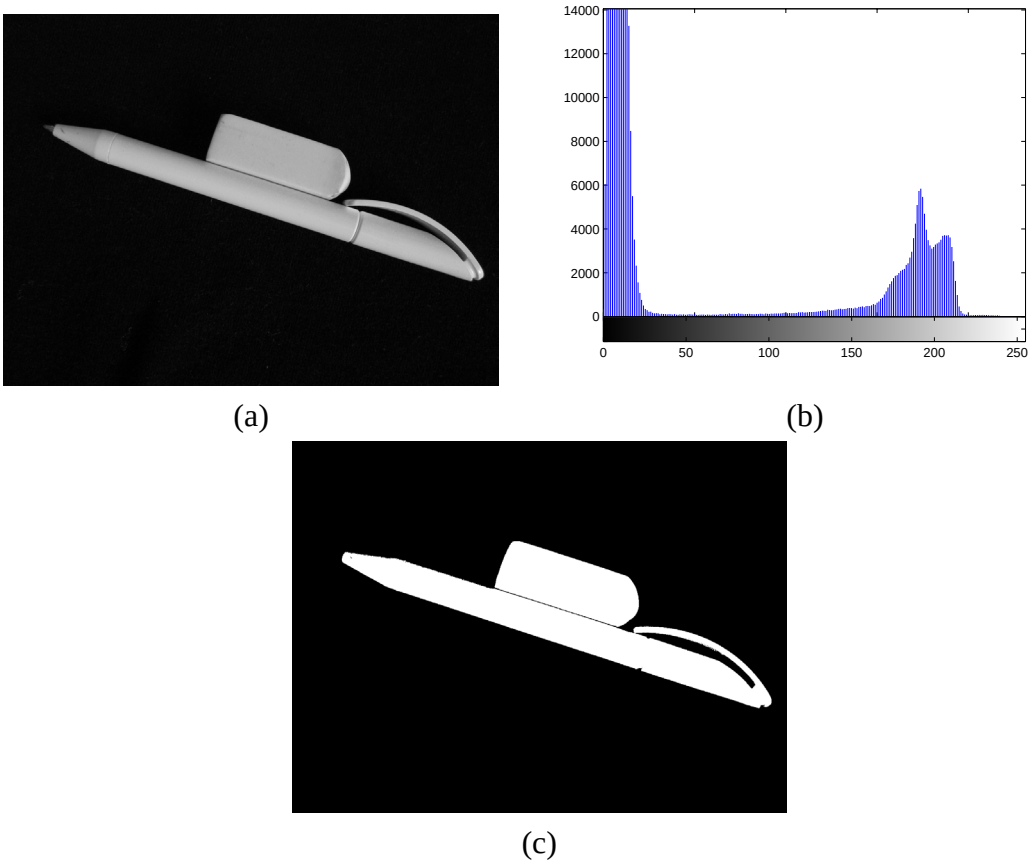

Figure 9: Example of thresholding: (a) Original gray-level image; (b) Histogram of the original image; (c) Result of thresholding the original image with threshold value 100.

Another example of global thresholding is shown in Figure 10. Figure 10(a) shows a gray-level image where the illumination is not uniform but increases linearly from left to right. Selecting a proper global threshold from the histogram shown in Figure 10(b) is now more difficult since the gray levels of the object and background are not separated as nicely as in the previous example. However, threshold value 84 is selected and the result of thresholding is shown in Figure 10(c). Global thresholding does not work very well now as over a half of the background pixels are incorrectly selected as belonging to the object. Even selecting a higher threshold (153) does not give a better result, since now dark object pixels are selected as belonging to the background.

As seen in the example above, global thresholding does not give adequately good results in images with non-uniform illumination, and thus local or adaptive thresholding or other image segmentation techniques should be used. The problem of non-uniform illumination will be discussed more precisely in Sub-section 4.2.2.

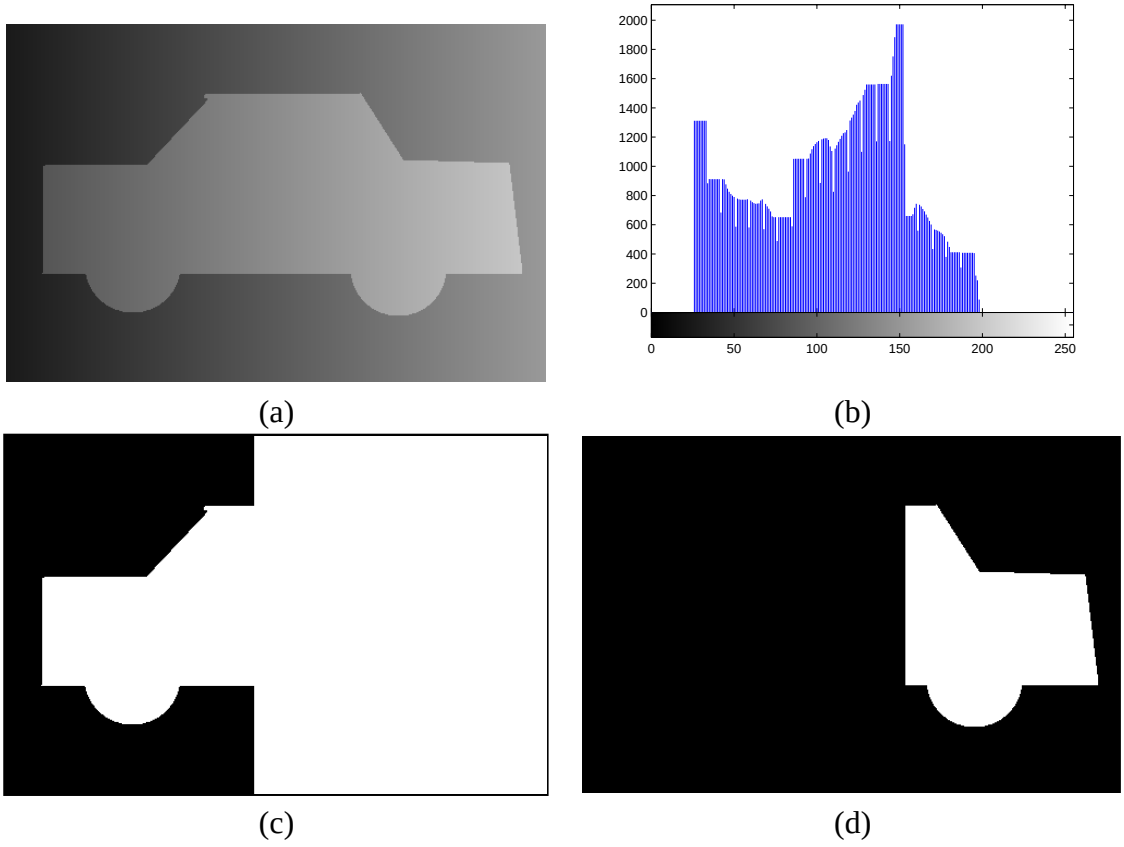

Figure 10: Example of thresholding an image with uneven illumination: (a) Original gray-level image; (b) Histogram of the original image; (c) Result of thresholding the original image with global threshold value 84; (d) Result of thresholding the original image with global threshold value 153.

2.2.2 Region growing

Region growing is an image segmentation method that groups pixels or subregions into larger regions based on some pre-defined criteria [11]. The method starts from a set of seed points. Regions where seed points belong are grown by appending those neighboring pixels that have properties similar to the areas grown so far. Neighboring pixels are specified by a connectivity rule, for example 4-connectivity or 8-connectivity, where the neighboring pixels are located in four or eight positions as shown in Figures 11(a) and (b), respectively.

Region growing is stopped when no more pixels satisfy the pre-defined criteria concerning for example gray-level, texture, or color. Additional criteria that increase the power of region growing method include size, region shape, or likeness between a candidate pixel

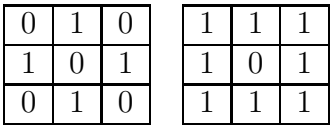

Figure 11: Examples of connectivities: (a) 4-connectivity; (b) 8-connectivity.

and the region grown so far.

Figure 12 shows an example of region growing. Figure 12(a) is an original gray-level image where a seed pixel is marked. The selected criterion for growing is that a neighboring pixel is added into the region if the absolute gray-level difference between the pixel and the region grown so far is less than or equal to 2. Neighboring pixels are selected by the 8-connectivity rule. In the first iteration round those pixels in the 8-neighborhood of the seed pixel are selected that satisfy the growing criterion. Since the neighboring pixels within range $[5, 9]$ are selected, only 4 is excluded, as shown in Figure 12(b). After the first growing round the mean value of the grown area is 7.6. During the second iteration round the pixels in the 8-neighborhood of any pixel in the grown area are checked. A neighboring pixel is selected if it is within range $[5.6, 9.6]$. Four neighboring pixels fit within the new range, and thus they are added into the area as shown in Figure 12(c). After the second iteration the mean gray-level value of the area is 8.0. In the third round all the neighboring pixels within range $[6, 10]$ are selected. Now only one neighboring pixel fit within the range and it is added into the area, as shown in Figure 12(d). After the third iteration round the mean gray-level value of the area is 8.2. Since none of neighboring pixels fit within range $[6.2, 10.2]$, the region growing is stopped.

2.2.3 Blob detection

The aim of blob detection is to find small regions that are either dimmer or brighter than their background in gray-level images. The basic principle of the blob detection method presented here is that a group of pixels in the original image are considered as a dark blob if all pixels in the group are darker than any of neighboring pixels surrounding the pixel group. Respectively, if all the pixels in the group are brighter than any pixel outside the group, the group is considered as a bright blob. Since maximum and minimum pixel values are searched in the method, even a single noise pixel (for example a bright noise pixel inside a dark blob) can make the method fail. Therefore the original image is filtered with a small median filter prior to blob detection.

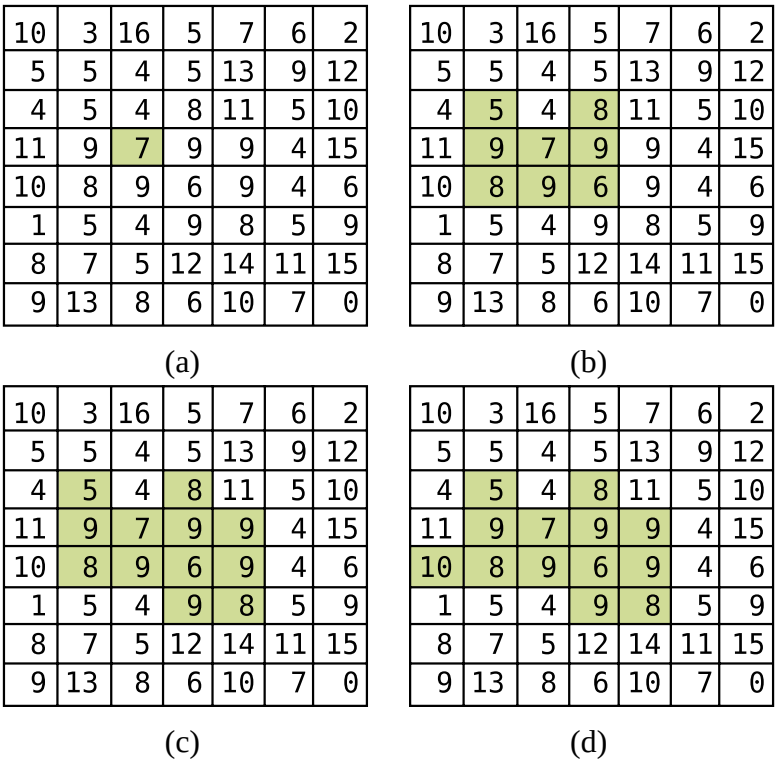

Figure 12: Example of region growing: (a) Original image where a seed pixel is marked; (b) Result after the first iteration; (c) Result after the second iteration; (d) Final result after the third iteration.

The blob detection method is implemented by using a circular filter that is calculated at every pixel in the image. There are three regions in the filter, as shown in Figure 13(a). Pixels belonging to a blob are considered to be inside the innermost circle and the neighboring pixels of the blob between the two outermost circles. The middle region is not used for calculations (demilitarized area), which makes it possible that a blob does not have to be circular. The rest of Figure 13 shows an example of filter masks needed in blob detection. Figure 13(b) is the inner mask for specifying the pixels of a blob. Figure 13(c) is the outer mask for specifying the neighborhood pixels of the blob. A combination of the two masks is shown in Figure 13(d). The filter finds blobs with a diameter between 5–11 pixels. When the size of the filter is increased, larger blobs are found. Thus, if the size of the blobs being searched may vary, several different mask sizes have to be used.

The method for finding dark blobs by using a single filter size is presented in Algorithm 1. The procedure may be repeated for different filter sizes if blobs of different sizes are searched. Filtering is performed for each pixel in an image, except for pixels near the edge of the image, as explained in Sub-section 2.1.1. First the maximum value of the image pixels specified by 1's in the inner mask is searched. Then the minimum value of

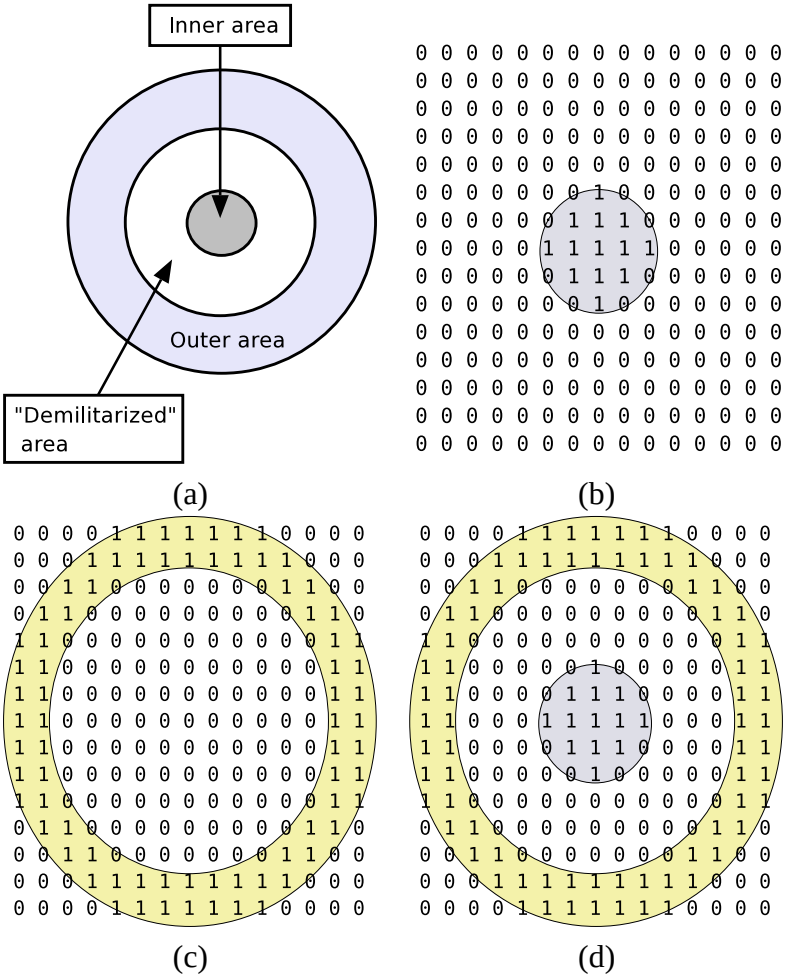

Figure 13: Masks for blob detection: (a) Mask naming; (b) Inner region mask specifying pixels belonging to a blob; (c) Outer region mask specifying neighboring pixels; (d) Combination of the inner and outer region masks.

the image pixels specified by the outer mask is searched. The minimum value is subtracted from the maximum value. If the result is less than a preset negative threshold, a dark blob is found, and thus, all 1's from the inner mask are marked into the result binary image. When all pixels have been processed, all found blobs have been marked in the result binary image.

Searching for bright blobs is processed similarly as in Algorithm 1, but now the minimum value is searched by using the inner mask and the maximum value by using the outer mask. The result appears by subtracting the maximum value from the minimum value. If the result is higher than a preset positive threshold, all the pixels inside the inner circle are marked to the result image.

Algorithm 1 Detection of dark blobs

1: **for all** pixels in the image **do**
2: Set the inner mask onto the image and select those image pixels that are marked with 1's in the mask.
3: Find the maximum value from the selected pixels.
4: Set the outer mask onto the image and select those image pixels that are marked with 1's in the mask.
5: Find the minimum value from the selected pixels.
6: Subtract the minimum value from the maximum value.
7: **if** the result is less than a preset (negative) threshold **then**
8: Add all 1's from the inner mask into the result mask.
9: **end if**
10: **end for**

An example of detecting bright blobs is shown in Figure 14. Figure 14(a) shows a gray-level image where illumination increases linearly from left to right. Due to the non-uniform illumination, global thresholding cannot be used for segmenting objects from the background. It is also visible in the figure that the objects do not have similar gray-levels, but objects in the right are brighter than objects in the left. However, it is possible to detect the blobs since they are brighter than their neighborhood. Figure 14(b) shows the result of using the described blob detection algorithm with the masks shown in Figure 13.

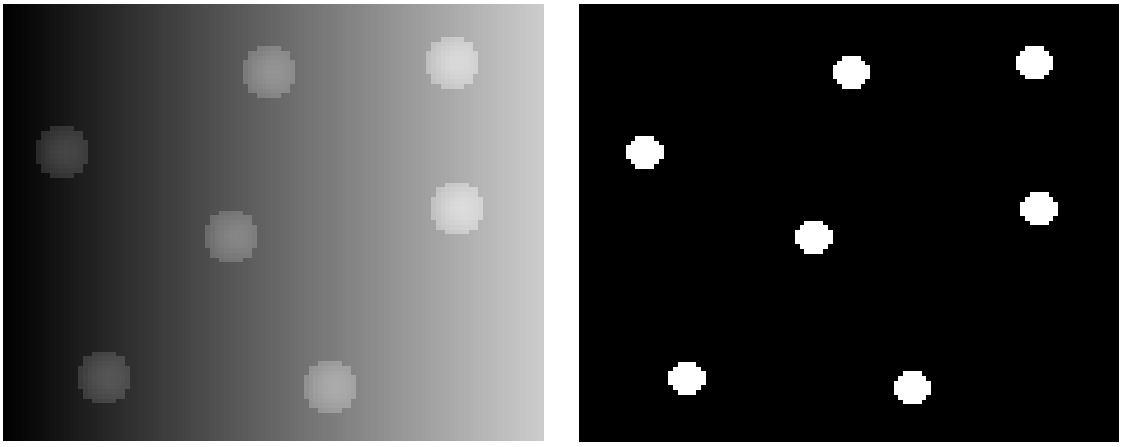

(a) (b)

Figure 14: Example of detecting bright blobs: (a) Gray-level image containing bright blobs; (b) Result of blob detection.

2.2.4 Color image segmentation

A color image usually consists of three channels that can be considered as separate gray-level images. Thus, segmentation of color images can be simplified to segmentation of gray-level images. The results of the separate segmentations are then combined into the final color segmentation result.

In the RGB color system, three color channels specify the fullness of red (R), green (G), and blue (B) color. It is possible to segment a color image in the RGB color system. Since the absolute color components, R, G, and B, contain intensity information in addition to color information, it is necessary to first convert the absolute color values to relative values R_r, G_r, and B_r that only specify the proportion of each color component but not its intensity. The conversion is done by dividing each absolute color component by the sum of the color components S as follows

$$R_r = \begin{cases} R/S & \text{if } S > 0 \\ 0 & \text{if } S = 0. \end{cases} \quad G_r = \begin{cases} G/S & \text{if } S > 0 \\ 0 & \text{if } S = 0. \end{cases} \quad B_r = \begin{cases} B/S & \text{if } S > 0 \\ 0 & \text{if } S = 0. \end{cases} \tag{4}$$

The sum S is simply calculated as

$$S = R + G + B. \tag{5}$$

After converting the color components of each pixel in the image into relative values, segmentation can be performed by using for example thresholding. Each of the color channels are thresholded as separate gray-scale images as described in Sub-section 2.2.1, and the binary result masks of thresholding are combined by applying the logical AND operation. Figures 15–17 show an example of searching for red objects. Figure 15(a) shows the original RGB color image. Figures 15(b–c) show color channels R, G, and B in the original image, respectively. Relative color component channels R_r, G_r, and B_r are shown in Figure 16. Let us assume that the task is to segment the thick red pen in the image. A selected object pixel belonging to the pen has relative color values 0.9275, 0.0048, and 0.0676 in the relative red, green, and blue color channels, respectively. Each of the relative color channels are separately thresholded. Pixels within ranges $[0.8, 1]$, $[0, 0.1]$, and $[0, 0.2]$ are selected in the relative red, green, and blue color channels, respectively. Finally the separate thresholding results are combined with the logical AND operation. The results of thresholding and the final segmentation result are shown in Figure 17.

There also exist more intuitive color systems, such as the HSI color system, that are

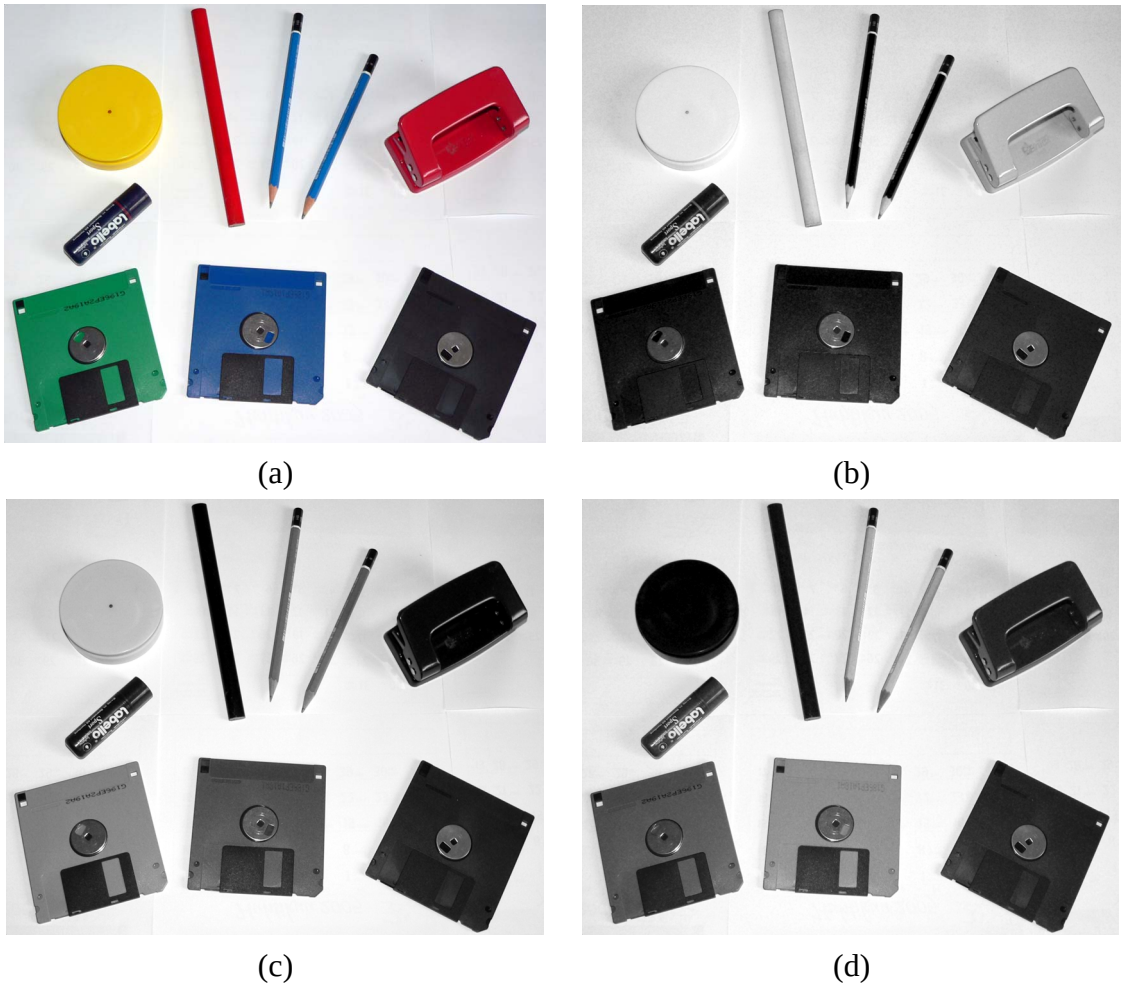

Figure 15: Color channels in the RGB image: (a) Original RGB image; (b) Red channel (R) in the original image; (c) Green channel (G) in the original image; (d) Blue channel (B) in the original image.

closer to the way a human being experiences colors. In the HSI color system hue (H) is a color attribute that describes a pure color (for example pure red) and saturation (S) specifies how much the pure color is diluted by white light. Intensity (I) is an achromatic parameter for brightness of objects. The hue component (H) for a certain pixel can be obtained from RGB values R, G, and B that are normalized to range $[0, 1]$ [11]

$$H = \begin{cases} \theta & \text{if } B \leq G \\ 360 - \theta & \text{if } B > G \end{cases} \tag{6}$$

with

$$\theta = \cos^{-1}\left\{ \frac{\frac{1}{2}\left[(R-G) + (R-B)\right]}{\left[(R-G)^2 + (R-B)(G-B)\right]^{1/2}} \right\}.$$

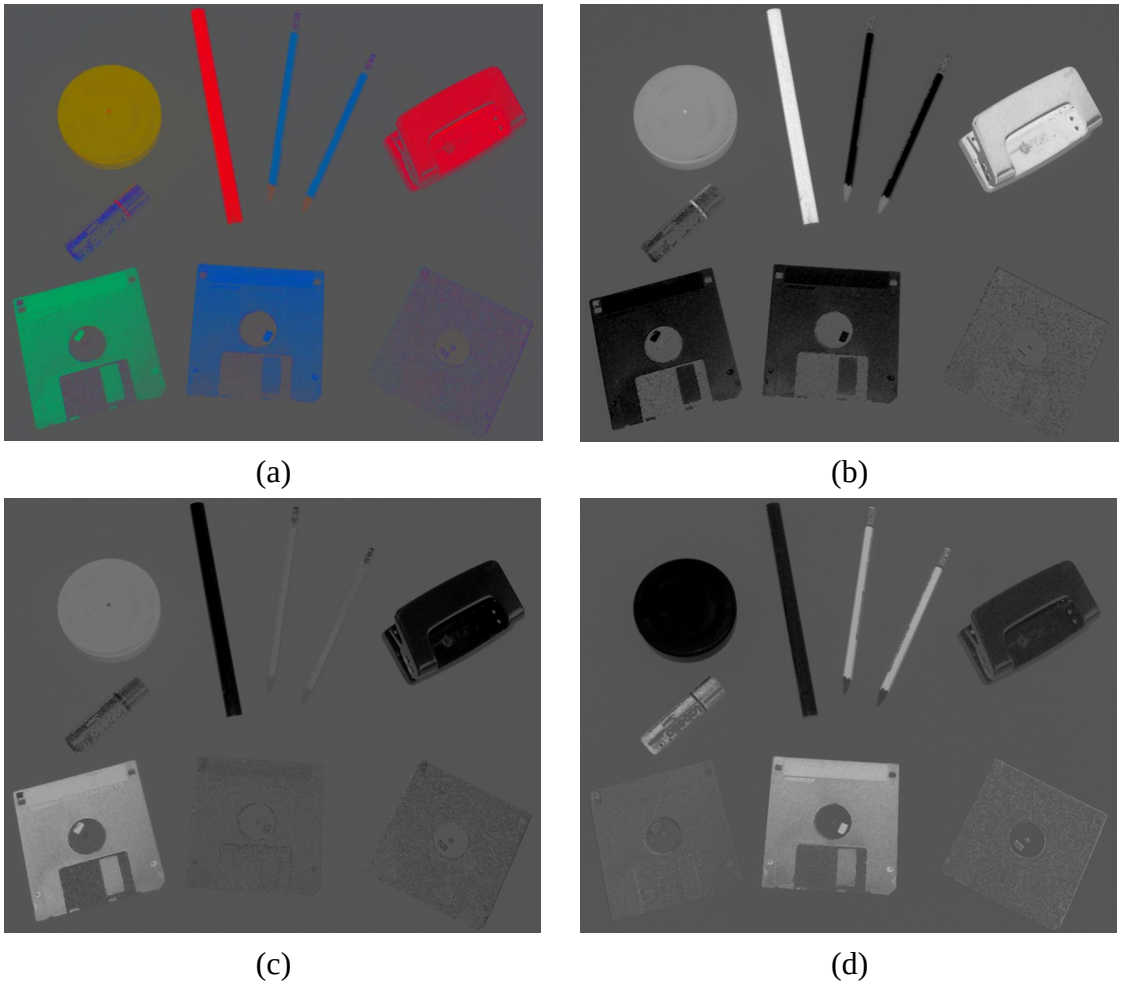

(a) (b) (c) (d)

Figure 16: Relative RGB color channels: (a) Relative RGB image; (b) Relative red channel (R_r); (c) Relative green channel (G_r); (d) Relative blue channel (B_r).

The saturation component (S) can be calculated from

$$S = 1 - \frac{3}{R+G+B}\left[min(R,G,B)\right]. \tag{7}$$

Finally, the intensity component (I) is given by

$$I = \frac{1}{3}(R+G+B). \tag{8}$$

The HSV color system is another color system similar to HSI. Also the HSV color system consists of hue (H), saturation (S), and intensity (V, value) components. Both the HSI and HSV color systems can be used in a similar way in color image segmentation. Figure 18 shows the H, S, and V channels got from the original RGB image shown in Fig-

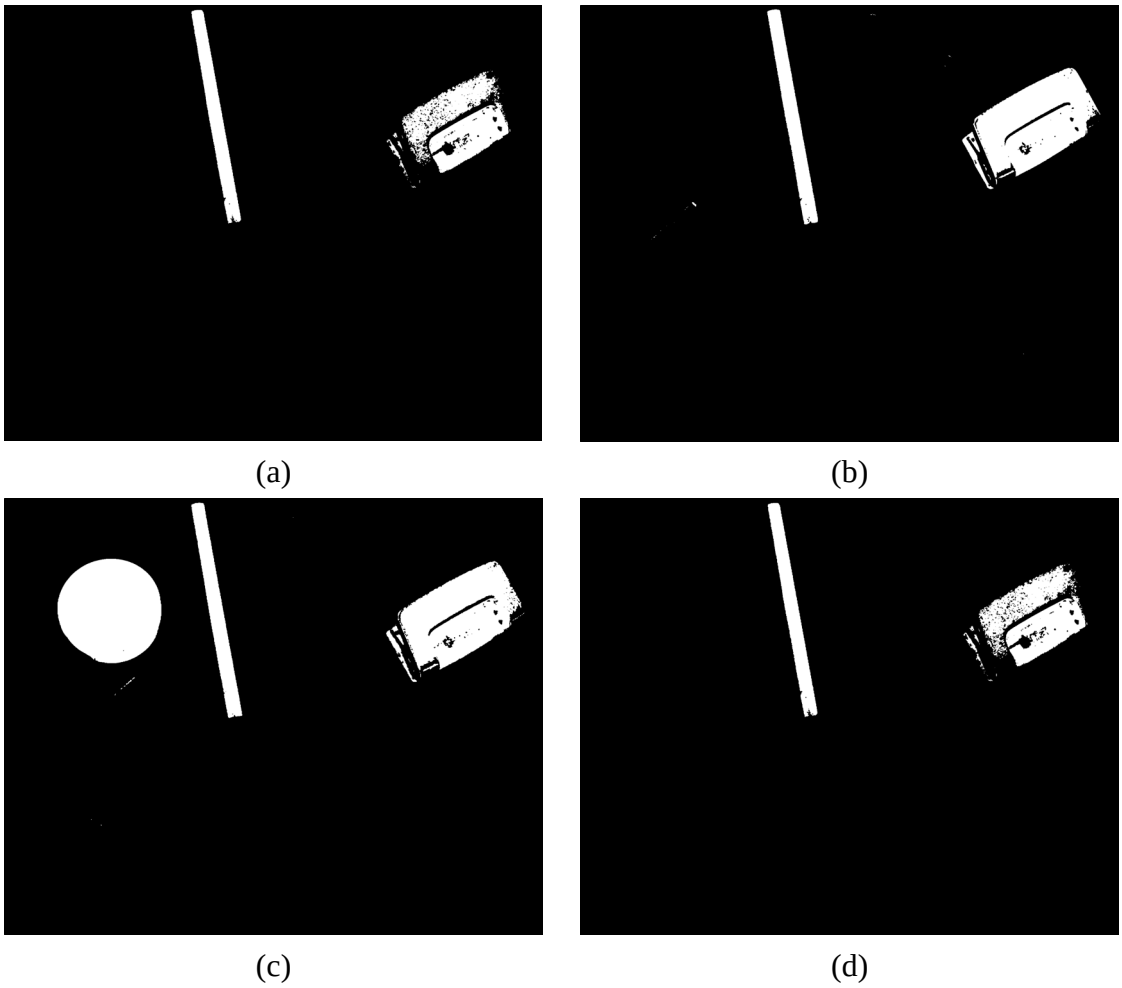

Figure 17: Example of the RGB color image segmentation: (a) Thresholded relative red channel (R_r) in Figure 16(a); (b) Thresholded relative green channel (G_r); (c) Thresholded relative blue channel (B_r); (d) Result of color segmentation, obtained by applying the logical AND operation to (a), (b), and (c).

ure 15(a). Since the hue channel describes the color, objects of red color can be detected by thresholding the hue channel shown in Figure 18(a). Hue values are not very reliable when the saturation is near to zero, and thus pixels having low saturation are excluded from the result. Red color in the HSV color system is within hue ranges $[0, 1/12]$ and $[11/12, 1]$, and thus these ranges are used when thresholding the hue channel. The saturation channel is thresholded with range $[1/2, 1]$. Figure 19 shows the thresholding results and the final color segmentation result obtained by applying the logical AND operation to the thresholding results.

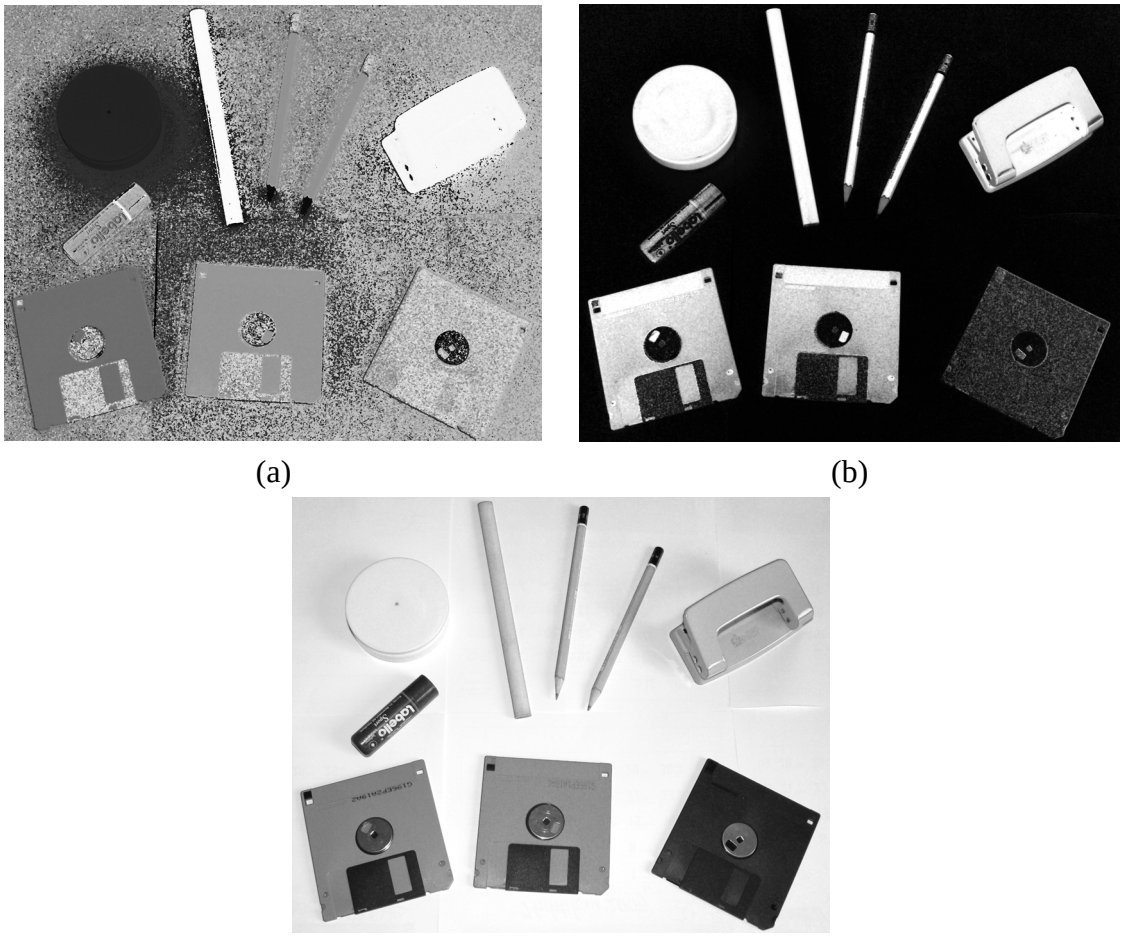

(a) (b)

(c)

Figure 18: Channels in the HSV image: (a) Hue channel (H) in an HSV image converted from Figure 15(a); (b) Saturation channel (S) in the HSV image; (c) Intensity (value) channel (V) in the HSV image.

2.3 Binary image processing

The result of color segmentation is usually a binary image where the objects are marked with 1's and the background with 0's. The binary image containing information of the objects can also be called a (binary) object mask since further image processing operations can be performed for the objects when omitting the background. Morphology can be used for processing binary images further, for example removing small objects or filling small holes. Binary object masks can also be used for measuring or calculating object features, such as size or mean intensity. The following sub-sections describe morphology and image measuring techniques.

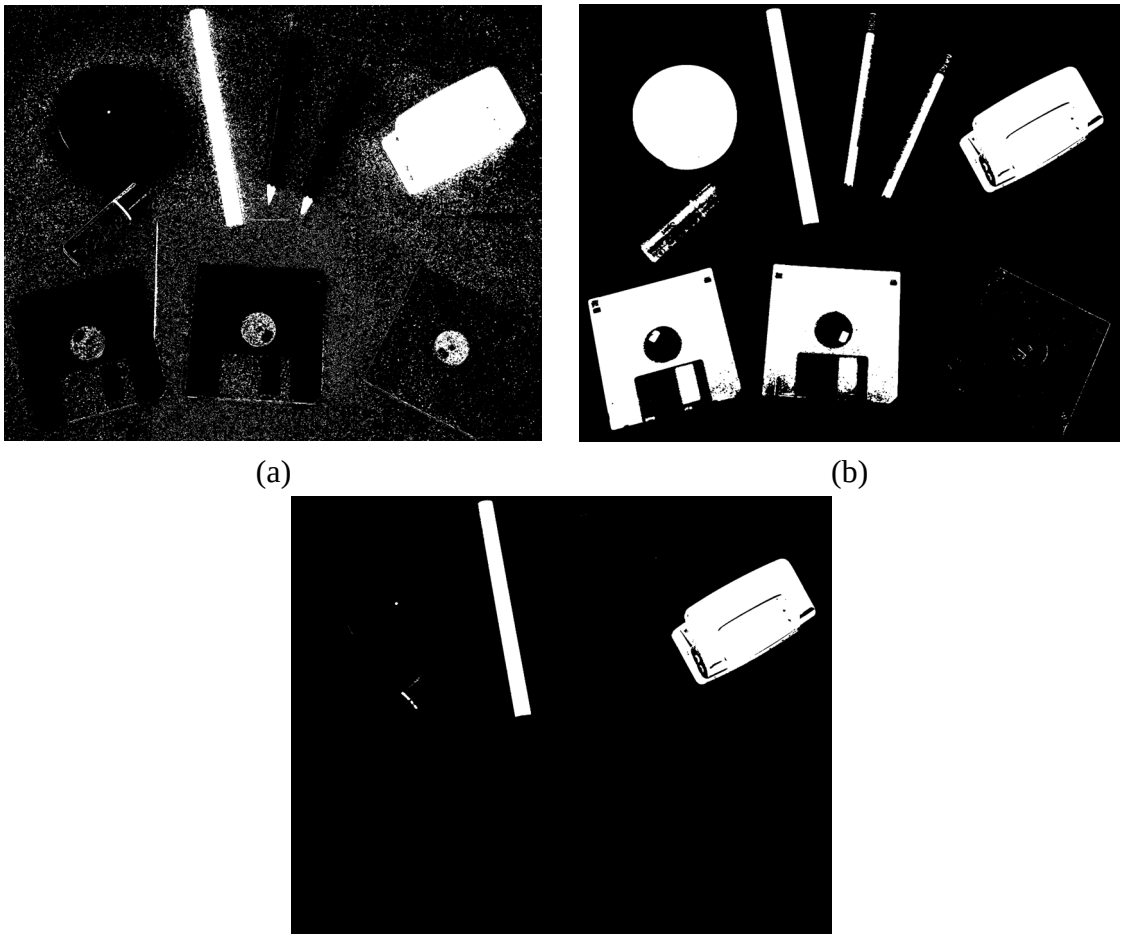

(a) (b)

(c)

Figure 19: Thresholding channels in the HSV image: (a) Multilevel thresholded hue channel, pixels within ranges $[0, 1/12]$ and $[11/12, 1]$ are selected; (b) Saturation channel thresholded with value 0.5; (c) Result of applying the logical AND operation to (a) and (b).

2.3.1 Morphological image processing

Mathematical morphology, usually performed for binary images, contains two fundamental operations: morphological dilation and erosion. Dilation expands and erosion shrinks objects marked in a binary image. Other morphological operations are for example morphological opening and closing which are based on dilation and erosion. Although morphology is often performed for binary images containing only 0's and 1's, morphology operations and algorithms can also be extended for gray-scale images. [11]

Morphological dilation and erosion are based on a certain mask called a structuring element. Figure 20 shows a few examples of different structuring elements. Both in dilation

and erosion, a mask or a structuring element, is moved on an image similarly as in filtering, which was explained in Sub-section 2.1.1.

0	1	0
1	1	1
0	1	0

(a)

1	1	1
1	1	1
1	1	1

(b)

1
1
1

(c)

Figure 20: Example structuring elements: (a) Disk-shaped structuring element; (b) Square-shaped structuring element; (c) Example asymmetric structuring element.

In dilation the center of the structuring element is checked on every pixel in the binary image. If the center pixel is marked with 1 in the image, all pixels specified by 1's in the structuring element are marked as 1's in the result image. Morphological dilation is expressed as $A \oplus B$, where A is the image being processed and B the structuring element. Figure 21 shows an example of dilation where the neighboring pixels near to an object are found by first dilating the binary image and then applying the logical exclusive-or (XOR) operation.

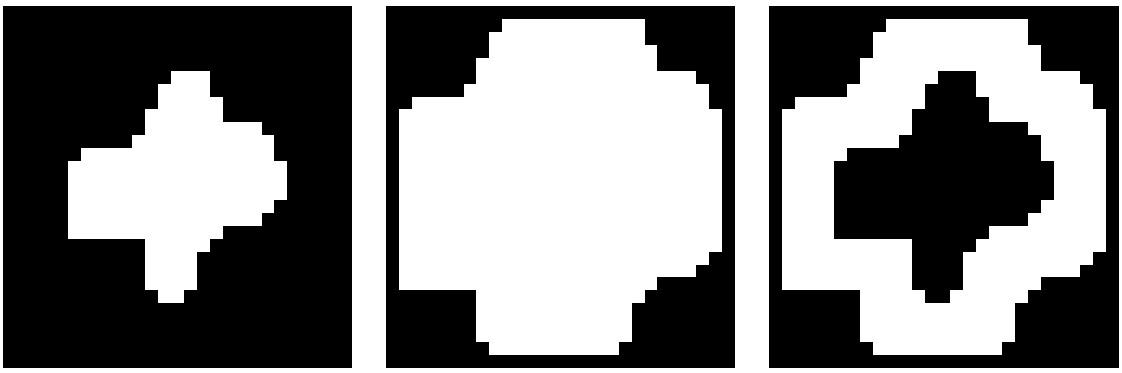

(a) (b) (c)

Figure 21: Example of morphological dilation: (a) Original binary image of size 27 x 28 pixels; (b) Result of dilating the binary image with a 9 x 9 square-shaped structuring element; (c) Neighboring pixels achieved by applying the logical XOR operation to (a) and (b).

Morphological erosion is the opposite operation to dilation. Also in erosion the center of the structuring element is checked on every pixel in the image, but if the center pixel is marked with 0 in the image, all pixels specified by 1's in the structuring element are marked with 0's in the result image. Morphological erosion is marked as $A \ominus B$, where A is the image being processed and B the structuring element. An example of erosion is shown in Figure 22 where the edge of an object is detected by first eroding the binary image and then applying the logical exclusive-or (XOR) operator.

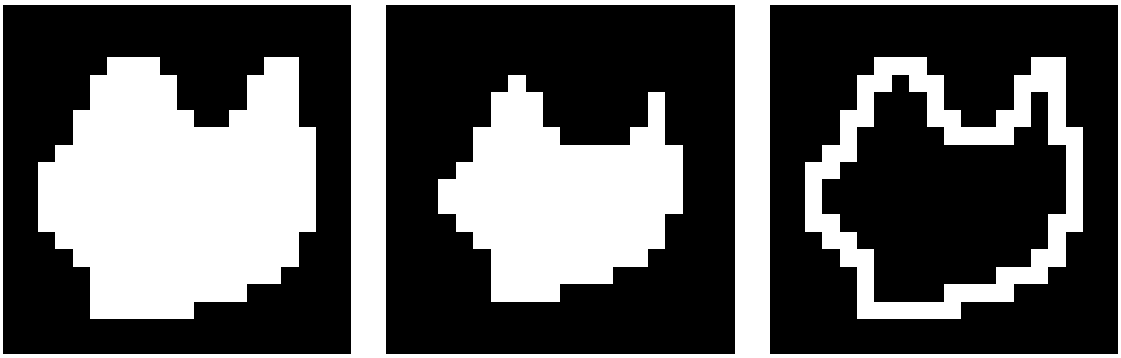

(a) (b) (c)

Figure 22: Example of morphological erosion: (a) Original binary image of size 20 x 20 pixels; (b) Result of dilating the binary image with a 3 x 3 square-shaped structuring element; (c) Edge pixels achieved by applying the logical XOR to (a) and (b).

In addition to dilation and erosion, morphological opening and closing are also important operations in morphology. Opening smoothes the shape of an object, eliminates thin protrusions, and removes narrow isthmuses [11]. Morphological opening, marked as $A \circ B$, can be performed by first eroding an image and then dilating the result of the erosion: $A \circ B = (A \ominus B) \oplus B$. An example of opening is shown in Figure 23. As the result of opening, the narrow part connecting the two rectangles is removed as well as the two small objects in the lower left corner of the image.

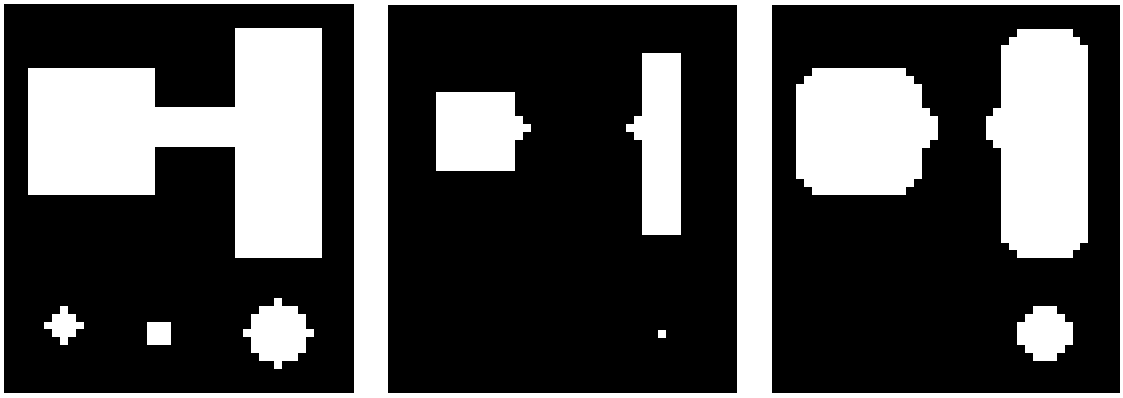

(a) (b) (c)

Figure 23: Example of morphological opening: (a) Original binary image of size 44 x 49 pixels; (b) Result of eroding the binary image with a disk-shaped structuring element; (c) Result of dilating (b) with the same structuring element.

Also closing smoothes the shape of an object, but as opposed to opening, it connects close parts of the object (or separate objects) together and fills small holes and gaps inside the object and in the object contour. Closing, marked as $A \bullet B$, can be performed by first

dilating an image and then eroding the result of the dilation: $A \bullet B = (A \oplus B) \ominus B$. An example of closing is shown in Figure 24, where a crack and a hole inside an object are filled and two objects close to each other are connected.

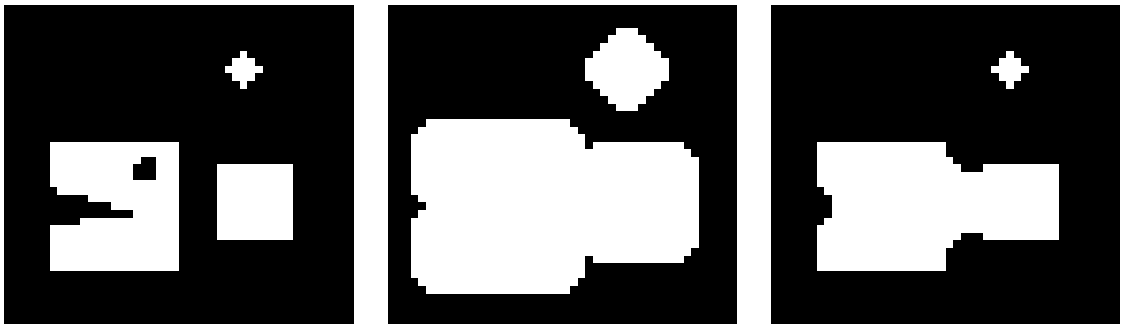

(a) (b) (c)

Figure 24: Example of morphological closing: (a) Original binary image of size 46 x 42 pixels; (b) Result of dilating the binary image with a disk-shaped structuring element; (c) Result of eroding (b) with the same structuring element.

2.3.2 Object descriptors

In machine vision applications it is often necessary to measure or calculate some properties of an object or several separate objects in an image. After image segmentation and marking each segmented object in a separate binary mask, the properties of objects can be measured or calculated.

The size of an object in pixels can be obtained by simply counting the number of pixels in the binary mask of the object. The average intensity value of an object can be calculated by summing those gray-level values in the original gray-scale image that are marked with 1's in the binary mask and dividing the sum by the object size in pixels. Sometimes it may be necessary to compare the intensity of an object to the intensity of the background surrounding the object. In this case, neighboring pixels can be detected by using morphological dilation as explained in Sub-section 2.3.1 above, and the mean intensity value of the neighboring pixels can be calculated similarly as the mean intensity value of the object.

There are some properties that describe the shape of an object, for example the object compactness and the ratio of the height and width of the object. Compactness means how tightly the pixels of an object are grouped together. A simple estimation for compactness can be calculated by dividing the number of object pixels by the number of pixels in a

rectangle that precisely fits the object. Figure 25(a) shows an example of an object with high compactness value. Another example is shown in Figure 25(b) where the object has a low compactness value as the pixels of the object are not very tightly grouped together, but there are a lot of background pixels inside the bounding rectangle. The ratio of the height and width of an object means how elongated the object is, and thus can be calculated by dividing the length of the shortest side of the rectangle by the length of the longest side. If the ratio is 1, it means that the object is square- or circle-like. If the ratio is lower than 1, it means that the object is oblong or oval. For example, the ratio 0.25 means that the longest side of the object is four times as long as the shortest side of the object. The example object in Figure 25(a) has a lower ratio than the example object in Figure 25(b).

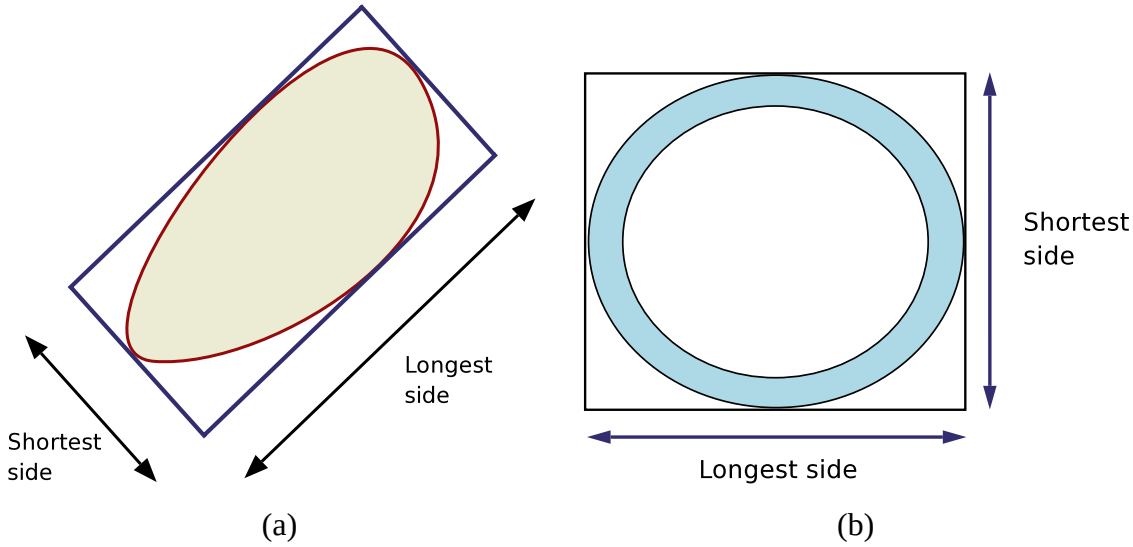

Figure 25: Object compactness and ratio: (a) Object with a high compactness and low ratio; (b) Object with a low compactness and high ratio.

2.4 Object classification

Since the aim of machine vision is to get information of objects in an image, the final step of the machine vision approach is usually object classification. In this step the candidate objects found in the image segmentation are classified as belonging to some of preset classes, or rejected if they do not belong to any of the classes. The class information of objects can be further used for example when deciding whether an object in the image is valid or whether it should be rejected due to deficiencies in quality.

There are several techniques for object classification. However, all classification methods need a knowledge database, in other words, prior knowledge of the problem domain. The

knowledge is usually provided by a domain expert, or experts, and is often called the ground truth. The domain expert labels objects in an image set and provides the label data and the image set to the team developing the classification methods. First, the provided image set and its ground truth are divided into two separate parts: a training image set and its ground truth to be used in the development of image segmentation and classification methods, and a testing image set with ground truth to be only used when testing the developed methods.

It is possible to develop a rule-based classifier. This means that some if-else rules are written that classify the candidate objects into different classes. The rules utilize information about objects measured from a segmented image. Although rule-based classifiers are easy to develop, they may be difficult to extend or apply to complicated problem domains.

The artificial neural network is a widely used classification technique, since it learns from input data and does not need any domain dependent rules to be written. The training of an artificial neural network is also simple since the network acts like a black box. This means that the weights inside the neural network are changed automatically during the training process. There are two types of neural networks that can be divided according to the type of the learning process: self-organizing maps (SOM) are unsupervised neural networks, whereas multilayer feedforward neural networks need supervised training. Multilayer feedforward neural networks are trained by giving train feature vectors and the corresponding classes to the network. If the output of the network does not correspond to the given class, the training process adjust the weights in the network [11]. The processing is continued until the network responds correctly to the given inputs. Self-organizing maps do not need prior knowledge of object classes but adjust network weights solely from input feature vectors used in training. The black-box feature of neural networks is also their disadvantage since it is not possible to formally prove that a neural network works in a correct way in every situation. Another disadvantage of neural networks is that they always classify objects into closest classes, but they do not express how reliably the objects belong to those classes.

Statistical classifiers are also used in classification applications. Statistical classifiers use probability functions for classifying candidate objects. The idea of a statistical classifier is to select the class where an object most probably belongs to. It is also possible to reject an object if it does not belong to any predefined class at least with a certain probability. Also several other classification techniques, such as structural classifiers, have been described in the literature, but it is not reasonable to discuss them in this master's thesis.

After the selected classifier is trained with the training image set and the corresponding ground truth, the classifier is tested against the test image set and the corresponding ground truth. For achieving reliable results, it is important that the test image set is adequately large, covers the domain where the developed system is aimed at and was not used when the methods were developed.

3 Diabetic retinopathy

Diabetic retinopathy is one of the complications caused by diabetes. As indicated by the name, diabetic retinopathy appears in the retina, which is the tissue responsible for vision in the eye. Since diabetic retinopathy causes changes in the eye, the disease may affect the vision. Diabetic retinopathy is the major cause of blindness among people of working age in developed countries [10]. It is more common in Type 1 than in Type 2 diabetes [10], [3], [13]. Half of Type 1 diabetics and one-fourth of Type 2 diabetics have some degree of retinopathy [14].

There are several risk factors that make it more possible to a diabetic to get diabetic retinopathy. The most important one is the duration of diabetes [3], since retinopathy does not appear suddenly, but it takes several years to develop [13]. Almost all of Type 1 and about half of Type 2 diabetics have some degree of retinopathy when they have had diabetes for 20 years [10], [13]. Another risk factor is the high level of blood glucose, which harms the blood vessels of the retina and disturbs their activity [13]. Other risk factors are high blood pressure, increased level of blood lipids, smoking, and pregnancy [3], [13], [14].

3.1 The structure of the eye

Before discussing the retina abnormalities caused by diabetic retinopathy, it may be useful to describe the structure of the eye, see Figures 26(a) and (b). Figure 26(a) is a cross section of the eye and Figure 26(b) is a fundus image taken from a healthy eye by a digital fundus camera. The fundus is the back of the eye and it can be seen through the pupil.

The cornea is normally only a half millimeter thick, but it is responsible for 70 % of the total focusing of the eye. The iris is a thin diaphragm located behind the cornea. It contains the pupil that dilates and constricts and thus controls the amount of light that reaches the retina through the lens. The task of the lens is to refract light. The lens can also change its shape to focus on objects at different distances. The focusing reflex is called accommodation and is done approximately 100,000 times in a day. The retina is the tissue where the image is projected. The retina contains light sensitive cells called cones and rods, which are responsible for daytime and night vision, respectively. The optic nerve transfers information of the projected image from the retina to the brain. The

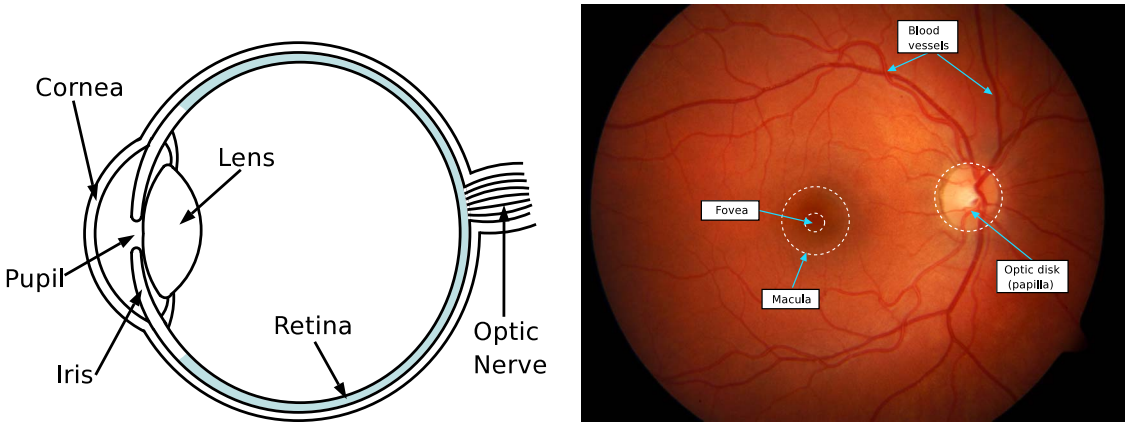

(a) (b)

Figure 26: Two representations of the eye: (a) Cross section of the eye; (b) Fundus image where the main parts of normal fundus are marked.

head of the optic nerve, called optic disk or papilla, does not contain receptors itself, and is thus the blind spot of the eye. [15], [16]

The macula is an oval spot near the center of the retina with a diameter of about 1.5 millimeters. The fovea is near the center of the macula and it contains packed cone cells. Due to high amount of light sensitive cells, the fovea is responsible for the most accurate vision. [16], [15]

3.2 Retina abnormalities in diabetic retinopathy

Diabetic retinopathy causes changes in the retina, which is the most important tissue of the eye [13]. An analogy between the eye and a camera could be presented to illustrate the importance of the retina: if an eye is considered as a camera, the retina is the film of the camera that catches an image from the target [13].

There may exist different kinds of abnormal lesions caused by diabetic retinopathy in a diabetic's eye. In this master's thesis the following lesion types are briefly described: microaneurysm, hard exudate, soft exudate, hemorrhage, and neovascularization. The features of different abnormal lesion types are shown in Table 1.

Microaneurysms are the earliest clinically detectable lesions [3], [13], [14], and thus it is important to have a method for detecting microaneurysms so that incipient retinopathy

Table 1: Abnormal lesion types in diabetic retinopathy.

Lesion type	Size	Color	Shape	Other
Microaneurysm	very small	dark red	round-shaped	
Hemorrhage	from small to large	dark red	dot or flame-shaped	
Hard exudate	from small to large	yellow	not regular	clear edges
Soft exudate	from small to medium	whitish	usually oval-shaped	blurry edges
Neovascularization	varies	red	varies	new blood vessels

may be noticed in its early stages. Microaneurysms are local distensions of retina capillary [14]. These outpouchings of capillary walls can be seen as tiny, round, red dots [3]. Figure 27(a) represents a microaneurysm marked with an arrow. The image is magnified, but it is still difficult to distinguish the microaneurysm due to its small size. When microaneurysms are coated with blood, they may be indistinguishable from dot hemorrhages [3]. Thus, the term "red small dot" is used in this master's thesis to cover both microaneurysms and small dot hemorrhages. Microaneurysms are not necessarily permanent changes, but they may first appear and then disappear during some period of time [13], [14].

Intra-retinal hemorrhages appear when capillaries or microaneurysms rupture [17] and some blood leaks out of the vessels. Hemorrhages can be seen as red, dot-blot or flame-shaped regions [3]. A fundus with multiple hemorrhages is shown in Figure 27(b).

Hard exudates are accumulated lipid formations leaked from weakened vessels [13], [17]. Hard exudate lesions are waxy and yellow with relatively clear edges [3]. Hard exudates often appear in clusters or rings [3]. Figure 27(c) represents a fundus where multiple hard exudates appear as bright lesions.

Soft exudates, also called cotton wool spots or micro-infarctions, appear when terminal retinal arterioles are obstructed [17], [13]. Soft exudates are small, whitish lesions with blurry edges [3]. Figure 27(d) shows a magnified fundus image where a soft exudate lesion is marked with an arrow. As seen in the figure, soft exudates are not usually as visible as hard exudates.

Extensive lack of oxygen caused by obstructions may lead to development of new blood vessels that are weak and can therefore easily tear [14], [18], [13], [17]. The disease where new blood vessels appear is called neovascularization. Neovascularization is the

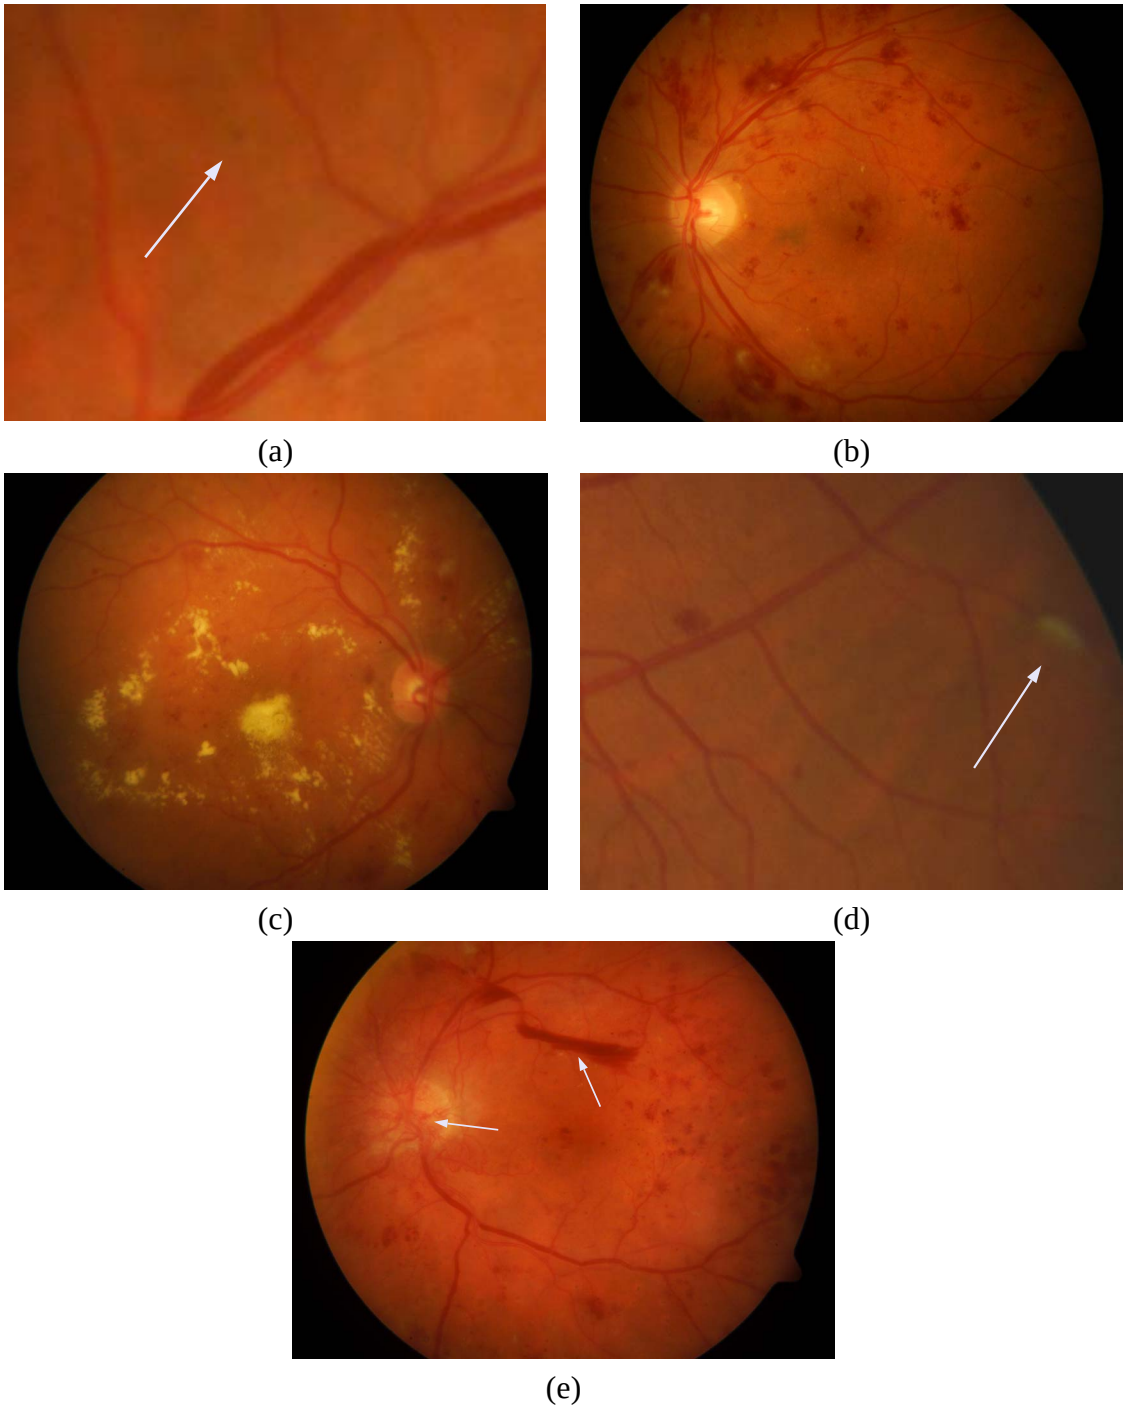

Figure 27: Abnormalities belonging to diabetic retinopathy: (a) Microaneurysm (small red lesion); (b) Hemorrhages (dark red lesions); (c) Hard exudates (yellow lesions); (d) Soft exudate (yellow/whitish lesion); (e) Neovascularization (new blood vessels).

most serious abnormality type in diabetic retinopathy since profuse bleeding may produce loss of vision [18]. Figure 27(e) shows a fundus with neovascularization. New small vessels have appeared on the optic disk that is marked with an arrow in the image. The other arrow shows a hemorrhage lesion where blood has leaked from a new blood vessel.

3.3 Screening for diabetic retinopathy

Sight-threatening retinopathy should be detected before it causes loss of vision. A proper treatment, like laser photo-coagulation (Figure 28) or glycemic control, may save the sight if the treatment is done early enough [19], [18]. Retinopathy is often asymptomatic, in other words, the patient is unaware of retinopathy until the eyes are routinely examined or until visual impairment is detected [19]. Therefore, diabetics' eyes should be regularly checked.

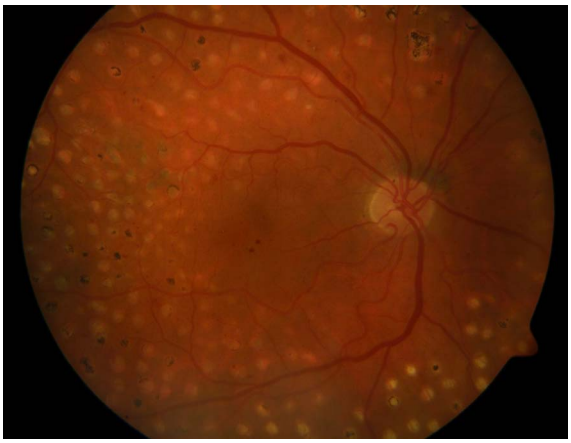

Figure 28: Fundus with laser scars after laser photo-coagulation.

It is unlike that significant retinopathy occurs before puberty, but then the diabetic's eyes should be examined at least every second year. If retinopathy appears, the eyes should be examined at least annually. [10], [14], [19], [18]

Fortunately there exist efficient ways for eye examination without making the patient feel uncomfortable during the operation. Since retinopathy lesions are visible in the fundus, the examination can be simply performed by viewing the funduses of diabetic's eyes using the natural way of the light: light is directed through the pupil to the retina and the fundus with its normal and abnormal parts can be observed from the reflected light.

The two most commonly used screening methods are ophthalmoscopy and retinal pho-

tography. Ophthalmoscopy can be direct or indirect, and it is an inexpensive, transportable, and widely available method. Retinal photography is more expensive than ophthalmoscopy, but it is more sensitive and produces permanent documentation. The main difference between the methods is that ophthalmoscopy should always be performed by a trained observer, whereas retinal photography can be manned by technical personnel, while the images are later reviewed by a medical specialist. [10], [19]

There exist several types of retinal photography. One way to register the light reflecting from the fundus is to use a color fundus camera. Two RGB color fundus images taken by a fundus camera are shown in Figure 29(a) and Figure 30(a). Alternatively, different filters and a gray-level camera can be used to catch wavelengths of interest. A very common method is to use a green filter for getting wavelengths of green color and register the filtered light using a gray-level camera. A gray-level fundus image created by using a green filter and a gray-level camera is shown in Figure 29(b). Note that Figures 29(a) and (b) are images of the same eye.

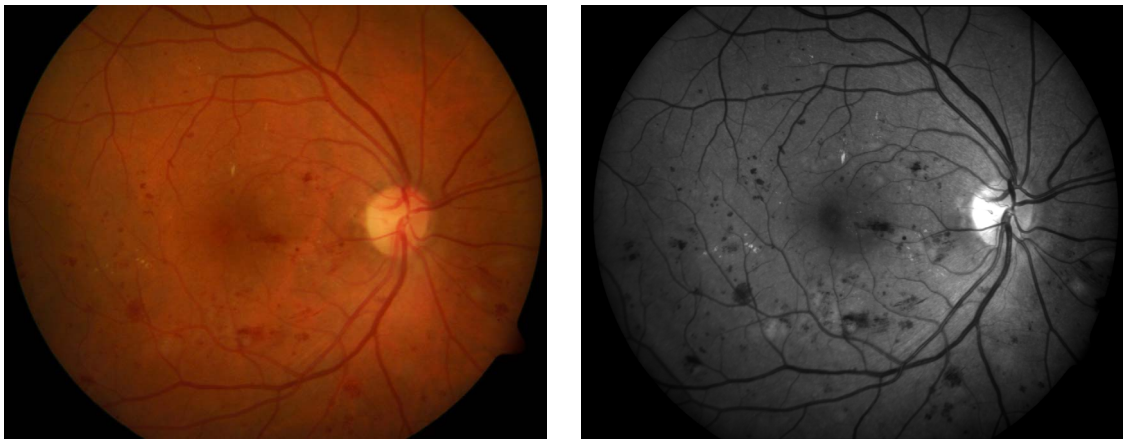

(a) (b)

Figure 29: RGB color fundus image and green filtered fundus image of the same eye: (a) Color fundus image; (b) Green filtered fundus image.

Another way for imaging the retina is to use fluorescein angiograms. In this method, a fluorescent liquid injection is given to the patient before the angiogram is taken. A fluorescein angiogram is shown in Figure 30(b), taken of the same eye as in Figure 30(a). In fluorescein angiograms, the parts of the retina where blood mostly flows are visible as bright whereas other parts of the retina appear dark. Thus, blood vessels and microaneurysms can be easily seen in fluorescein angiograms since they appear brighter than other parts of the retina. Hemorrhage lesions contain blood that does not flow, and thus they are visible as dark regions in fluorescein angiograms. However, fluorescein angiograms have a

disadvantage that makes them inconvenient: the need of an injection.

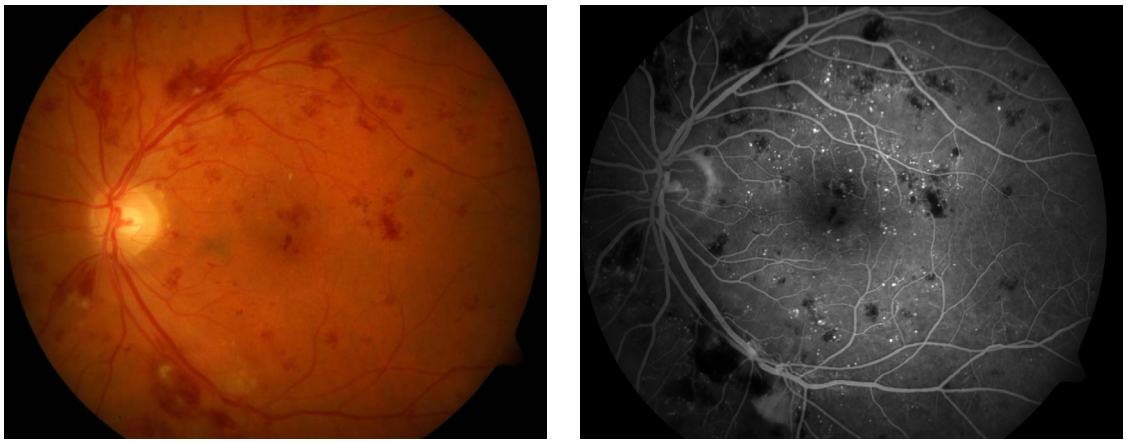

(a) (b)

Figure 30: RGB color fundus image and fluorescein angiogram of the same eye: (a) Color fundus image; (b) Fluorescein angiogram.

At the moment the screening of diabetic retinopathy is performed by trained medical experts. Diabetes UK guidelines recommend that any procedure used for screening sight-threatening diabetic retinopathy should have at least 80% sensitivity and 95% specificity [20]. Sensitivity means the percentage of abnormal funduses classified as abnormal by the procedure. Specificity means the percentage of normal funduses classified as normal by the procedure. The higher the sensitivity and specificity values, the better the procedure. Sensitivity and specificity values can be calculated as follows [21]

$$sensitivity = \frac{T_P}{T_P + F_N} \tag{9}$$

$$specificity = \frac{T_N}{T_N + F_P} \tag{10}$$

where T_P, T_N, F_P, and F_N mean true positives, true negatives, false positives, and false negatives, respectively. A screened fundus is considered as a true positive if the fundus is really abnormal and if the screening procedure also classified it as abnormal. Similarly, a true negative means that the fundus is really normal and the procedure also classified it as normal. A false positive means that the fundus is really normal, but the procedure classified it as abnormal. A false negative means that the procedure classified the screened fundus as normal, but it really is abnormal.

Since fundus images are nowadays in the digital format, it is possible to create a computer-

based system that automatically detects abnormal lesions from fundus images. An automatic screening system would save the workload of well-paid ophthalmologists, letting hospitals and eye clinics to use their resources in other important tasks. It could also be possible to screen more people and more often with the help of an automatic screening system, since it would be more inexpensive than screening by humans. One advantage of the automatic screening system is that it is deterministic, in other words, it always classifies funduses in the similar way. There will always exist differences in the backgrounds and education of human screeners, which causes dispersion in their diagnose making. Also a single human expert may make different diagnoses in different screening times due to human factors, such as tiredness or sickness.

A computer-based screening system does not have to be perfect to be used in screening. It is better to use a computer-based screening system for classifying only clearly normal funduses as normal, whereas abnormal and obscure funduses are delivered to a human expert for further classification. However, the computer-based screening system reduces the workload of the human expert, since in the screening most of the funduses are normal, and only a few funduses have retinopathy.

However, if the first phase of the screening is performed automatically by a computer, it is important to have as high sensitivity as possible so that no retinopathic funduses are missed by the computer. Also in the case of the combined computer and human screening the above mentioned minimum sensitivity and specificity values should be achieved.

In addition to evaluating the screening methods by calculating the sensitivity and specificity values for the whole fundus images, a computer-based screening method can also be evaluated pixel-wise. In the pixel-wise evaluation the sensitivity and specificity values are calculated for pixels of fundus images rather than whole images. The pixel-wise evaluation is not very commonly used in the literature. Thus also the results in this thesis concern whole fundus images.

3.4 Machine vision in related diabetic retinopathy studies

There exist some literature about the detection of retinopathy lesions by computer-based methods. Most of the relevant studies concern exudate (especially hard exudate) detection from color fundus images. However, there are only a few studies on hemorrhage or microaneurysm detection from color fundus images. In addition, most microaneurysm

studies seem to concern only fluorescein angiograms. This sub-section lists the lesion detection methods and their results described in the literature.

Usher et al. [20] propose methods for detecting microaneurysms, hemorrhages, and exudates in color fundus images. The methods utilize local contrast enhancement and color standardization methods when pre-processing color fundus images. Since the color of the fundus may vary in different people, a color standardization technique is used to fit the colors of a fundus image to the colors of a reference image. Abnormality detection is handled by a region growing algorithm, an adaptive intensity thresholding, and an edge enhancement operator. The classification is done by an artificial neural network. The study is based on a large number of images (from 1273 patients) for training and testing and reports relatively good results: the sensitivity was 94.8% and the specificity 52.8% when considering the existence or absence of retinopathy rather than existence or absence of single lesion types. However, two images per patient (one from each eye) were used in the study when deciding whether the patient had retinopathy. Using two images per patient may have improved the results compared to other studies where only one image per patient has been used.

Sinthanayothin et al. [22] have used thresholding and neural networks for detecting microaneurysms and hemorrhages in color images. They have used a region growing algorithm for detecting hard exudates. They report 80.21% sensitivity and 70.66% specificity for detecting retinopathy. Totally 767 fundus images were used, but it is not mentioned whether the same images were used when developing and testing the algorithms.

Gardner et al. [23] divide red free fundus images (images where wavelengths of red color are omitted) into sub-blocks and use artificial neural networks for classifying the sub-blocks. They report 88.4% sensitivity and 83.5% specificity for detecting retinopathy as a whole (all lesion types together), 73.8% sensitivity and 73.8% specificity for hemorrhages, and 93.1% sensitivity and 93.1% specificity for exudates. They have used totally 301 color fundus images. Since most of the test images were also used when the system was trained, the reported sensitivity and specificity values are not very reliable.

Ege et al. [24] have compared how well three statistical classifiers, Bayesian, Mahalanobis, and K-Nearest Neighbor classifiers, are able to classify microaneurysms, hemorrhages, hard exudates, and soft exudates. They have used the region growing algorithm for getting candidate lesions from color fundus images, and the statistical classifiers for classifying candidate lesions into four classes. They report Mahalanobis classifier as having the best results, but unfortunately they list only the sensitivity values for different classes

but not any specificity values. Thus, the results of their research are not comparable with the other ones presented here.

Wang et al. [25] have used statistical classification (MDD classifier) for detecting exudates in 154 color fundus images. They report 100% sensitivity and 70% specificity, but they do not mention whether the same images were used in training and testing. Goh et al. [26] have used Bayes' theorem in their method to find exudates. They report 100% sensitivity and 73.9% specificity with totally 333 color fundus images. Only 23 images contained exudates and the test images were the same as the train images. Since the used fundus images were not divided into separate train and test sets and since there existed only a few exudate images, the results are not very reliable. Li and Chutatape [27] combine the region growing method and edge detection. They report 100% sensitivity and 71% specificity, but the system was tested with only 35 images and they do not mention, whether the same images were used in training and testing. In [28] and [29] the thresholding is used, but no results are reported. Hsu et al. [30] have used dynamic clustering method and domain knowledge for identifying true hard exudates. They used totally 543 color images in their project and report 100% sensitivity and 74% specificity. Osareh et al. [31], [32] have used fuzzy c-means clustering (FCM) for getting candidate hard exudate lesions. They compare how efficiently a neural network and support vector machine (SVM) can classify the candidate hard exudates. They report that the neural network provided better results (sensitivity 95.0% and specificity 88.9%) than the support vector machine (sensitivity 87.5% and specificity 92.5%), but they recommend the support vector machine, due to its advantages. Their image set was small, 40 images contained hard exudate and 27 images were normal.

Raman et al. [33] have concentrated on microaneurysm detection in color fundus images by using top-hat transformation and classifying candidate lesions according to lesion shape and color. They report 90% sensitivity and 87.5% specificity, but their image set contained only 18 images and the same images were used both in training and testing. There exist quite a lot of publications concerning microaneurysm detection in fluorescein angiograms. Cree et al. [34] report 82% sensitivity and 84% specificity. They used 68 fluorescein angiograms and the testing was done using a separate set of 20 fluorescein angiograms. Also, for example in [35], [36], [37], and [38] microaneurysm detection using fluorescein angiograms has been researched, but unfortunately neither sensitivity nor specificity values have been reported.

In the literature, a wide scale of techniques are described for detecting retinopathy lesions in fundus images or fluorescein angiograms. Sensitivity and specificity values are com-

monly used for evaluating how accurately a method classifies fundus images as abnormal or normal. However, it is impossible to compare the results of different studies due to several weaknesses in testing environments. Some studies do not report any numerical results or they only report either the sensitivity or specificity value. Reporting only one value does not tell anything of the goodness of a method since it is always possible to adjust the method so that it achieves either 100% sensitivity or 100% specificity: a method that classifies all images (both normal and abnormal) as abnormal has 100% sensitivity and another method classifying all images as normal has 100% specificity. Sensitivity and specificity values depend on each other; increasing the sensitivity of a method usually decreases its specificity. Thus, a compromise between sensitivity and specificity should usually be done. It is very important to report both sensitivity and specificity values so that it is possible to compare the results of different methods. There are also other problems that make it impossible to compare the current studies in the literature: all studies use their own fundus image databases, most of databases have an inadequate number of images, and/or the images are not divided into separate training and testing image sets. The only way to reliably compare different methods and studies is to use a common fundus image database where a large number of fundus images are divided into separate training and testing sets.

4 Fundus image acquisition and pre-processing

The RGB fundus images used in this research were provided by the Department of Ophthalmology at the University of Kuopio, Finland. The images were taken by a Zeiss Fundus Camera FF 450plus IR with a 50-degree field of view (Figure 31). The resolution of the fundus images is 1500 x 1152 pixels and the precision is 24 bits per pixel. In the following sub-sections it is discussed how the lesion data was collected and how the acquired fundus images were pre-processed in this study.

Figure 31: Zeiss Fundus Camera.

4.1 Computer aided data collection

Since the author of this thesis and the supervisors did not have medical knowledge or experience with fundus images when the research started, there was a need for a procedure for collecting and delivering information of abnormal lesions in fundus images. A Matlab-based tool with a graphical user interface was developed by the author in the beginning of the research (Figure 32). With the help of this tool, an ophthalmologist was able to efficiently and quickly mark abnormal lesions in fundus images. The tool was developed to be universal, and thus it can also be used for marking any similar objects in images of

other problem domains.

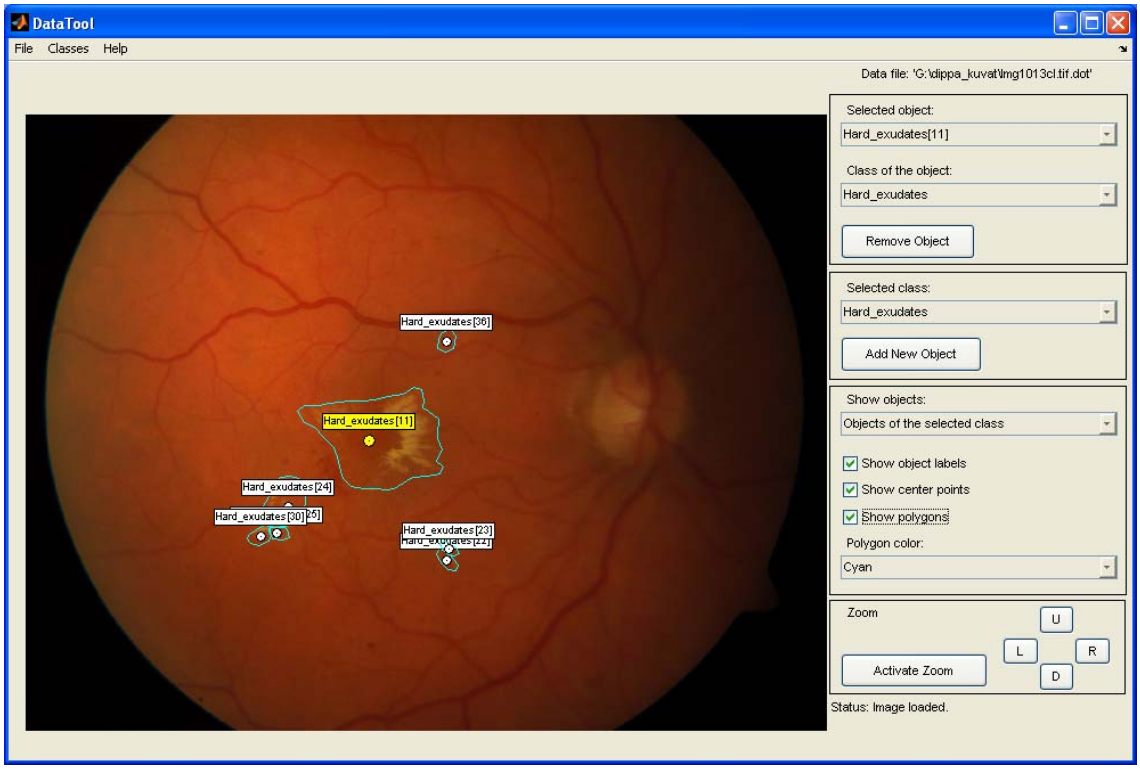

Figure 32: Tool for collecting lesion data.

It is possible to define a preset list of object classes that are available when the tool is started. The user can add more classes during the marking. In addition, the user can open a fundus image and zoom and scroll it. When marking a lesion (Figure 33), or any object, the user can select the class of the object and mark its center. It is also possible to mark a polygon-shaped line specifying the border of the object. The possibility to mark the edge of objects was considered as an important feature, since it is often difficult for a non-medical person to distinguish the scope and shape of a lesion.

In the tool the user is allowed to change the class of a marked object or remove the object if needed. The user can select whether all marked objects, only the selected object, or objects belonging to the selected class are shown. The user can also select how much information of the objects is shown (center points, labels, and/or edge lines).

When the objects have been marked, the user can save the marked data so that the image with the marked objects can be opened and displayed later. The tool saves the data into a text file and no changes are made into the image. The marked center points of objects, classes, and possible edge lines are saved into the file. The file can be opened and modified

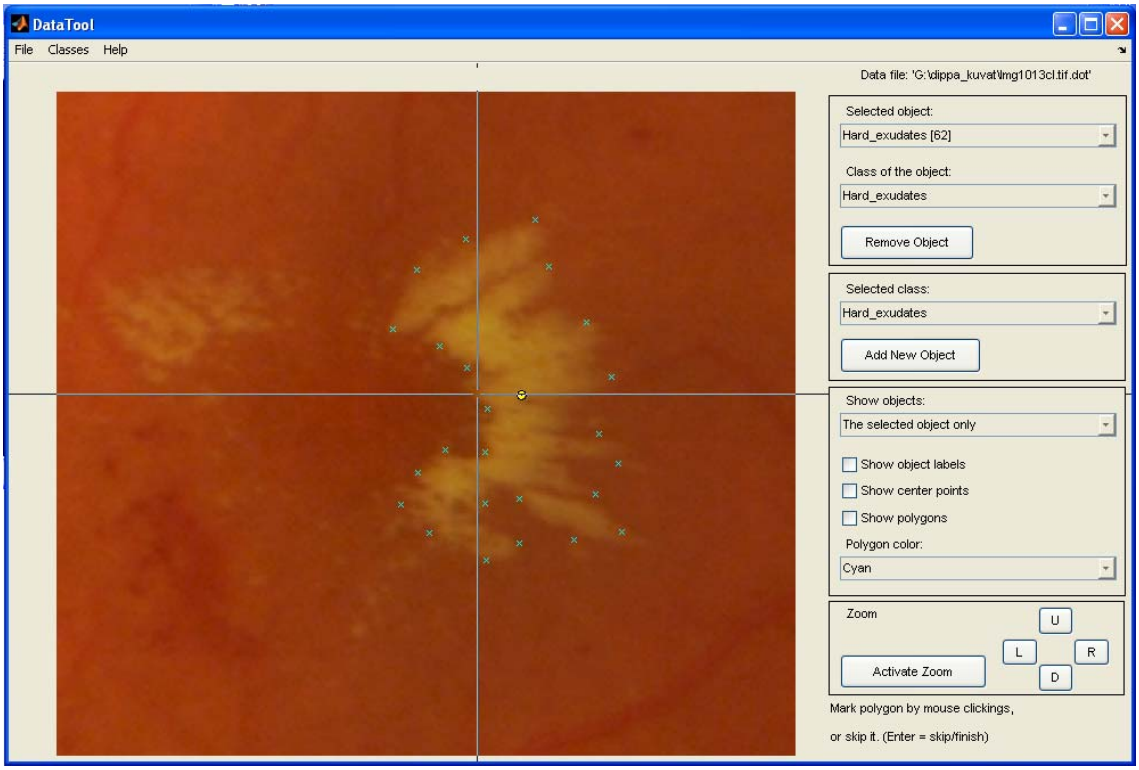

Figure 33: Marking lesions with the tool.

later by the tool or used directly in algorithms.

The created data file is a text file where each object is described in one line. The format of a line is as follows

x_c y_c [x_1,x_2,...,x_n,] [y_1,y_2,...,y_n,] $Class\_name$

where x_c and y_c are the x- and y-coordinates of the center point of an object, x_1, x_2, ..., x_n and y_1, y_2, ..., y_n are the x- and y-coordinates of the polygon line specifying the edge of the object, respectively. The polygon coordinates are optional, but the brackets are mandatory. The separator of the polygon coordinates is a comma, and it is mandatory also after the last coordinate. Finally, $Class\_name$ represents the name of the class where the object belongs. Note that empty spaces are not allowed in the class name since empty space is used for separating different fields of the line in the text file.

An example line of a data file could be as follows:

```
77 58 [42,59,98,93,72,42,] [44,95,72,39,25,44,] Hard_exudates
```

The center point is at $(77, 58)$, where 77 and 58 are the x- and y-coordinates, respectively. The x-coordinates of the polygon-shaped edge line are 42, 59, 98, 93, 72, and 42. The y-coordinates of the edge line are 44, 95, 72, 39, 25, 44. The class of the object is $Hard\_exudates$.

If a polygon-shaped edge line is not marked, a line of a data file could be, for example, as follows:

```
1043 533 [] [] Hemorrhages
```

4.2 Image pre-processing

Before abnormal lesions can be searched from an acquired fundus image, the image has to be pre-processed to ensure adequate level of success in the abnormality detection. Figure 34 shows the five parts of fundus image processing described in this thesis. Four of the parts belong to pre-processing: color space conversion, fundus region detection, illumination equalization, and poor image quality detection.

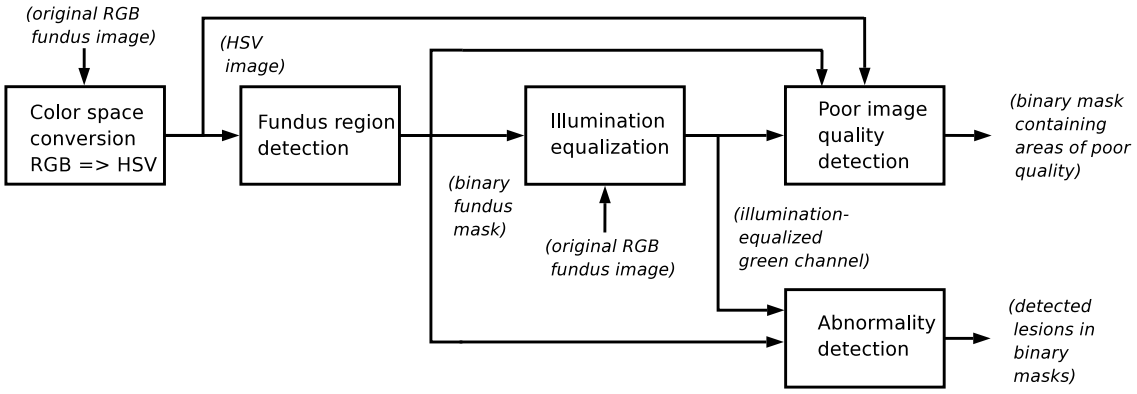

Figure 34: Parts of abnormal lesion detection with their inputs and outputs.

In the first part the acquired RGB color fundus image is converted to the HSV color system. The second part, fundus region detection, searches the region of the fundus from the HSV image. The result of the second part is a binary image where the found fundus is marked with 1's and the background of the fundus with 0's. The created fundus mask is given as an input to the third part, illumination equalization. This part also takes the original RGB image as an input and equalizes uneven illumination in the image. The result of the part is a gray-level image representing the green channel of the original

image where the illumination has been equalized. The fourth pre-processing part detects regions of the image having poor image quality. The poor image quality detection part takes the converted HSV image, the binary fundus mask, and the illumination-equalized green channel as inputs. The part produces a binary mask as output, where regions having poor image quality are marked with 1's. Finally, the results of pre-processing can be used in the abnormality detection, which is the fifth part of the fundus image processing. The abnormality detection part takes the binary fundus mask and the illumination-equalized green channel as inputs and produces separate binary masks for each of the lesion types searched for in the part. Note that the abnormality detection part searches for lesions in the whole fundus so that lesion information inside poor image quality regions is also available if needed.

The first part, color space conversion from RGB to HSV was already presented in Section 2. The three other pre-processing parts are described in the following sub-sections. The abnormality detection part is presented in Section 5. The parameter names of the used methods are briefly mentioned in the text in *italics*. All parameters and their default values are presented in Appendix 1.

4.2.1 Fundus region detection

A fundus image consists of a circular fundus and a dark background surrounding the fundus. It is important to separate the fundus from its background so that the further processing is only performed for the fundus and not interfered by pixels belonging to the background. In this sub-section a method for creating a binary fundus mask prior to lesion detection is described. In a fundus mask, pixels belonging to the fundus are marked with 1's and the background of the fundus with 0's. With the help of the fundus mask a lesion detection algorithm can process only the pixels of the fundus and omit the background pixels.

The fundus can be easily separated from the background when the original fundus image has been converted from the RGB color system to a color system where a separate channel is used to represent the intensity values of the image. As explained in Sub-section 2.2.4, HSV is such a color system. Figure 35(a) shows an original RGB color fundus image. The intensity channel (V) in the converted HSV image is shown in Figure 35(b).

The intensity channel can be thresholded by a low threshold since the background pix-

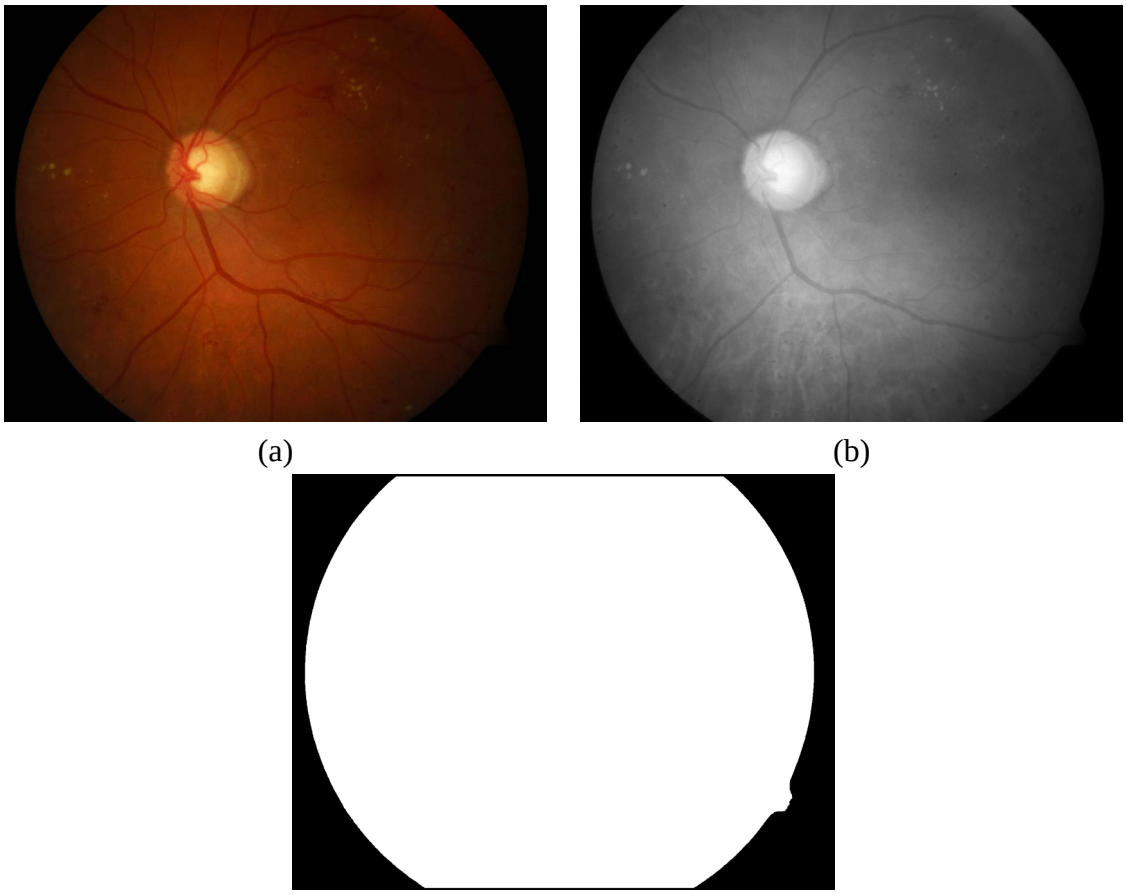

Figure 35: Creating a fundus mask: (a) Original RGB fundus image; (b) Intensity channel (V) in the HSV color system; (c) Fundus mask created by thresholding the intensity channel, removing fundus edge pixels, and filling holes inside the fundus region.

els are typically significantly darker than the fundus pixels. In some cases the intensity difference between the fundus and the background is not very clear due to inadequate illumination near to edge of the fundus. Thus, some dark regions of the fundus may be considered as background after thresholding. However, very dark pixels near to the fundus edge could be excluded from further processing in any case, since it is not possible to detect lesions from regions where illumination has been almost zero.

After thresholding, a small median filter is performed to remove single noise pixels from the created fundus mask. The edge of a fundus tends to be very noisy, and thus, the edge pixels are removed by morphological erosion with a small structuring element (*clearEdgeThickness*). The remaining holes in the mask can be filled by processing the mask line by a line and filling the missing pixels between the first and last fundus pixels in each row. A created fundus mask is shown in Figure 35(c).

4.2.2 Equalization of uneven illumination

A two-dimensional image $f(x, y)$ may be characterized by two components: amount of source illumination $i(x, y)$ and reflectance $r(x, y)$, which specifies how much of the illumination is reflected by the objects in the scene. The two functions combine as a product to form the image $f(x, y)$ [11]

$$f(x, y) \quad = \quad i(x, y)r(x, y). \tag{11}$$

The reflectance component, $r(x, y)$, specifies how much of the illumination coming to an object is reflected and how much is absorbed. If illumination $i(x, y)$ is known, the reflectance can be easily calculated from

$$r(x, y) \quad = \quad \frac{f(x, y)}{i(x, y)}. \tag{12}$$

Fundus images are characterized by uneven illumination: the center region of a fundus image is usually highly illuminated while the illumination decreases closer to the edge of the fundus. In other words, objects of the fundus (lesions and blood vessels) are differently illuminated in different locations of the image due to the non-uniform illumination. If the source illumination is known, the reflectance component can be easily calculated from Equation 12. Unfortunately, the source illumination is not typically known in the case of fundus images, but it depends on many unknown factors, for example ophthalmologic defects, misalignment of fundus camera, or eye movement [28]. However, an estimation of source illumination should be available to estimate the reflectance component of the objects. If the estimation $e(x, y)$ is available for $i(x, y)$, estimation of the reflectance component is

$$r(x, y) \quad \approx \quad \frac{f(x, y)}{e(x, y)}. \tag{13}$$

There are several methods in the literature for solving the problem of non-uniform illumination in fundus images. In [39], [40], [27], and [41] fundus images are divided into smaller sub-blocks, and the blocks are processed instead of a whole image. The idea of the method is that the illumination variation in a small block is lower than in the whole image. However, there are some problems with the method. For example adjacent pixels near to the border of two sub-blocks may have very different values in a result image depending on which block the pixels belong to. The problem is due to independently processed sub-blocks where distinct blocks do not take other blocks into account. In [20],

[22], [42], [31], [32], [43], and [44] the local contrast enhancement method is used for equalizing uneven illumination in the intensity channel of fundus images. In [28], [29], [30], and [24] a large mean filter, large median filter, or both are used for estimating the fundus background. 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 the background estimation. Also some other methods have been presented in the literature. For example, in [25] the intensity values in dark regions have been increased.

All the illumination equalization methods found in the literature were tested for the present study. Each method has its own advantages and disadvantages, but no technique was found that would completely solve the problem of uneven illumination. Since the results of the methods were not completely satisfactory, another method for illumination equalization had to be selected.

An image processing handbook [12] mentions a very simple, but effective method for illumination equalization in fundus images. Surprisingly the method had not been used in any studies found in the literature. In the proposed equalization method the source illumination of an RGB color fundus image is simply estimated by the red channel of the image, or more compactly, $e(x, y) = f_{red}(x, y)$, where $f_{red}(x, y)$ is the red channel of the RGB image. If we wish to equalize uneven source illumination in the green channel $f_{green}(x, y)$ and get its reflectance component $r_{green}(x, y)$, Equation 13 can be written as

$$r_{green}(x, y) \approx \frac{f_{green}(x, y)}{f_{red}(x, y)}. \tag{14}$$

The selected illumination equalization method has some difficulties in the regions of fundus images where the illumination has been inadequate. The problem turns regions that used to be shadowed in the original green channel into bright regions in the result image. The problem causes errors when searching for bright regions in the image. Thus, large median filtering was selected to supplement the selected illumination equalization method. This combination of two methods seemed to work very well.

Figure 36 shows an illumination equalization example using the described technique. The original image is the same as in Figure 35. The green channel of the original RGB image is shown in Figure 36(a). The illumination-equalized green channel obtained as a result of the proposed technique is shown in Figure 36(b). Note that there exist much noise in the edge and background of the fundus. The fundus mask shown in Figure 35(c) has been used to exclude background and fundus edge pixels. The result of masking the illumination-equalized green channel by the fundus mask is shown in Figure 36(c).

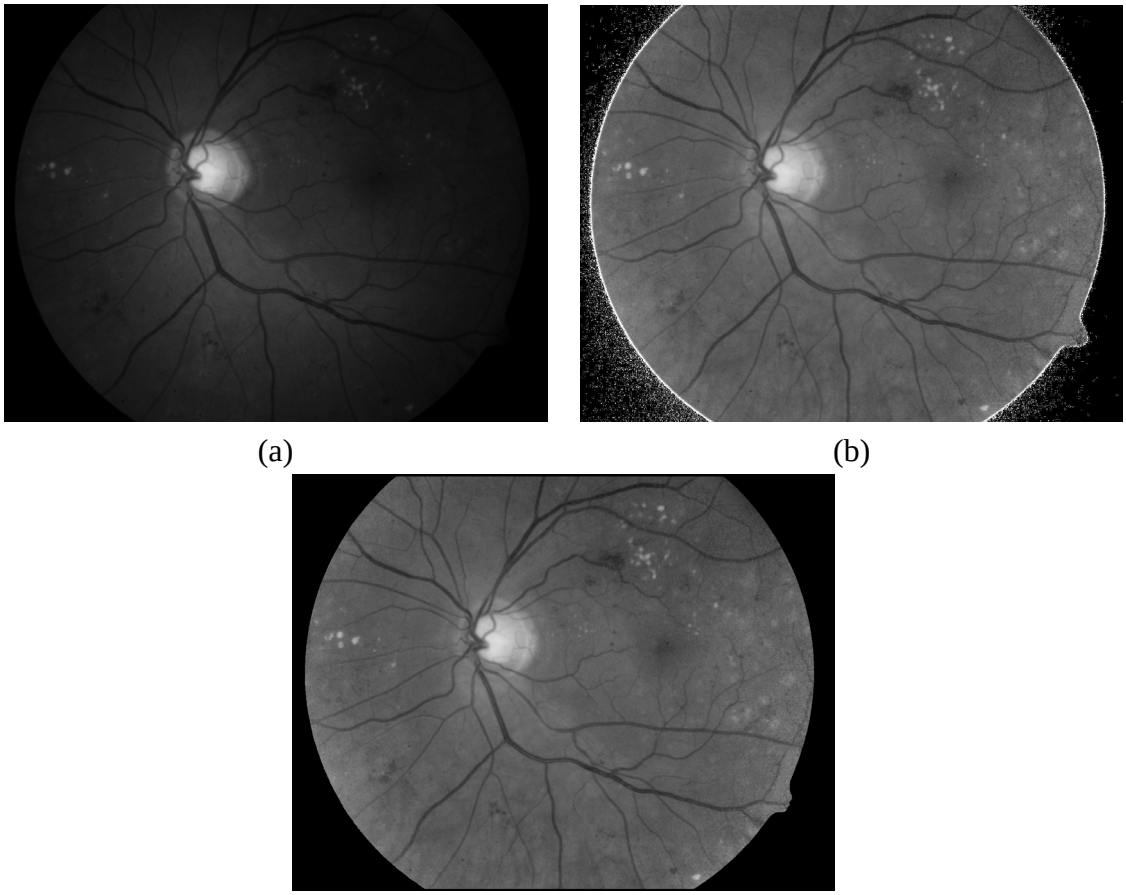

Figure 36: Example of illumination equalization: (a) Green channel of the original image shown in Figure 35(a); (b) Green channel of the original image where illumination has been equalized; (c) Result of masking the illumination-equalized green channel by the fundus mask shown in Figure 35(c).

Another illumination equalization example is shown in Figure 37. The illumination has been inadequate in the right-hand-side of the original RGB fundus image shown in Figure 37(a). As a result of poor illumination, it is difficult even for human eye to see if there are abnormal lesions in the shadowed regions. The same problem is visible in the green channel of the image, shown in Figure 37(b). The proposed illumination equalization technique performs well, and the result masked by a fundus mask can be seen in Figure 37(c). However, the lack of illumination has affected to the result of equalization; there exist noise in regions which were shadowed in the original image. The noise is more clearly visible in a magnified image, shown in Figure 37(d). Despite of noise appearing in the result image it is now possible to see that there are no serious lesions in the rightmost regions of the fundus.

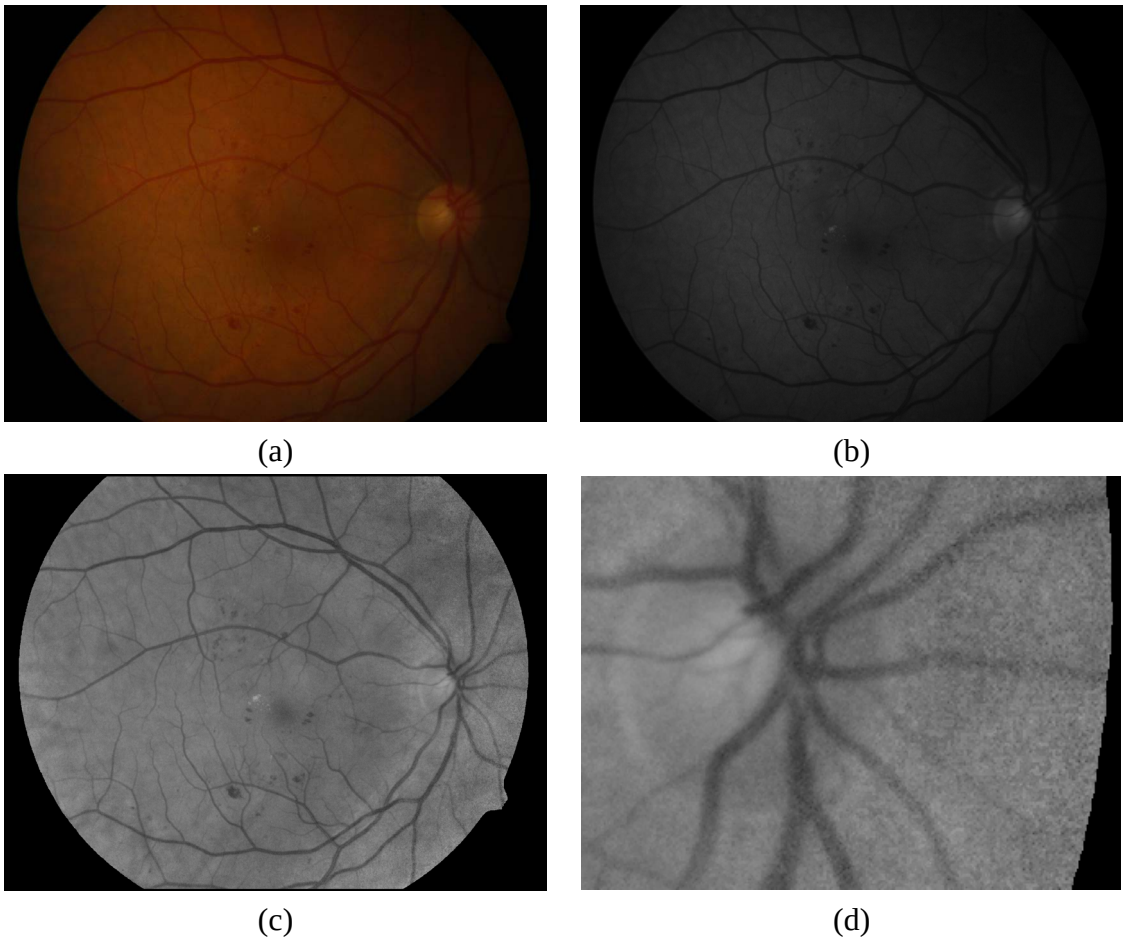

Figure 37: Example of illumination equalization in a poorly illuminated fundus image: (a) Original RGB fundus image; (b) Green channel of the original image; (c) Illumination-equalized green channel masked by a fundus mask; (d) Magnified area of the illumination-equalized green channel showing noise.

4.2.3 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, abnormal lesions are not searched from poor quality regions. However, in the second solution there is a risk that a fundus image having retinopathy is diagnosed to be normal if lesions only

exist in poor image quality regions. If the exclusive way is used, it is important that the algorithms provide information on what parts of the fundus image 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 project 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 existence or absence of diabetic retinopathy is made according to regions of good image quality and the user of the algorithms is informed of regions having inadequate image quality.

The detection of poor image quality regions is performed in five steps in this research. First regions having a high amount of noise are detected. Also regions where the fundus color is distorted are searched for. Then these two region groups are combined. Since the result of the combination may contain small gaps, the next step is to fill those holes. Finally artificial rectangular-shaped contours appearing in the result of the process are smoothed due to the assumption that distorted regions do not form rectangular-shaped blocks.

Noisy pixels can be found from a fundus image by applying a small median filter to the illumination-equalized green channel and by subtracting the median filtered image from the original one. The result of subtraction is thresholded, in other words, if the difference between a median filtered pixel and the original pixel is higher than a preset threshold (*randomNoiseSensitivity*), the pixel is considered as noise. After the noisy pixels have been marked into a binary mask, the mask is divided into small sub-blocks (*randomNoiseBlockWidth, randomNoiseBlockHeight*). If a sub-block contains more noise pixels than a preset threshold (*randomNoiseLimit*), the whole sub-block is marked to be noise. Figure 38 shows an example where noise pixels are detected from the illumination-equalized green channel presented in Figure 36(c). Figure 38(a) visualizes the found noise pixels superimposed on the fundus mask. Also the sub-blocks are marked in the figure. Sub-blocks containing more noise pixels than the preset limit are marked into a binary mask

and the result is shown in Figure 38(b).

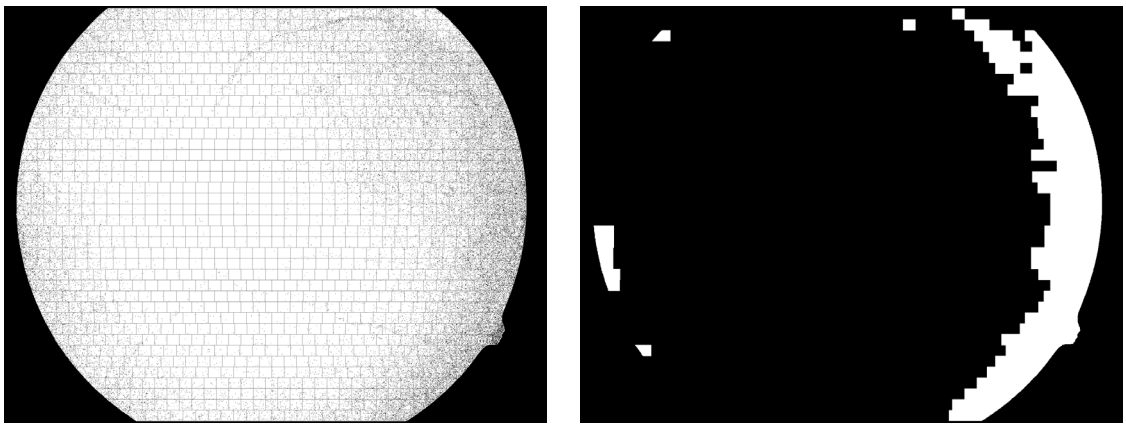

(a) (b)

Figure 38: Detecting noise pixels from the illumination-equalized green channel shown in Figure 36(c): (a) Fundus mask divided into sub-blocks where noise pixels are marked with black dots; (b) Sub-blocks having more noise pixels than a preset threshold.

In addition to noise, inadequate illumination may cause regions where the color of the fundus is distorted. It was found from the training image set that regions of distorted color have high hue values (H) and relatively low intensity values (V) in the HSV color system. Thus, regions having distorted color can be found by first dividing the hue channel by the intensity channel and then thresholding the result with a preset threshold (*badHueChannelThreshold*). Since pixels of distorted color seem to form large continuous regions, small regions are out of interest. Thus, small regions are removed by filtering the thresholded channel by a medium size median filter. Note that the median filter can also be applied to binary images; in binary images the median value is always either 0 or 1. Figure 39(a) shows the hue channel in an HSV image converted from the original RGB image shown in Figure 35(a). The hue channel divided by the intensity channel and thresholded by the preset threshold is shown in Figure 39(b). The result of median filtering the thresholding result is shown in Figure 39(c). The figure represents regions of the fundus where the color is distorted due to poor illumination conditions.

Regions containing a high amount of noise or distorted fundus color are combined into a single binary mask, or poor image quality mask, as shown in Figure 40(a). Since noise pixel detection was performed in artificial blocks, the block-shape is still visible in the poor image quality mask. It is assumed that noise is not distributed in block-shaped regions in images, but rather in smooth regions. Thus, tight corners in the poor image quality mask are rounded. The rounding of corners is performed by morphology operations. Morphology opening is first used for rounding the protruding parts and then the

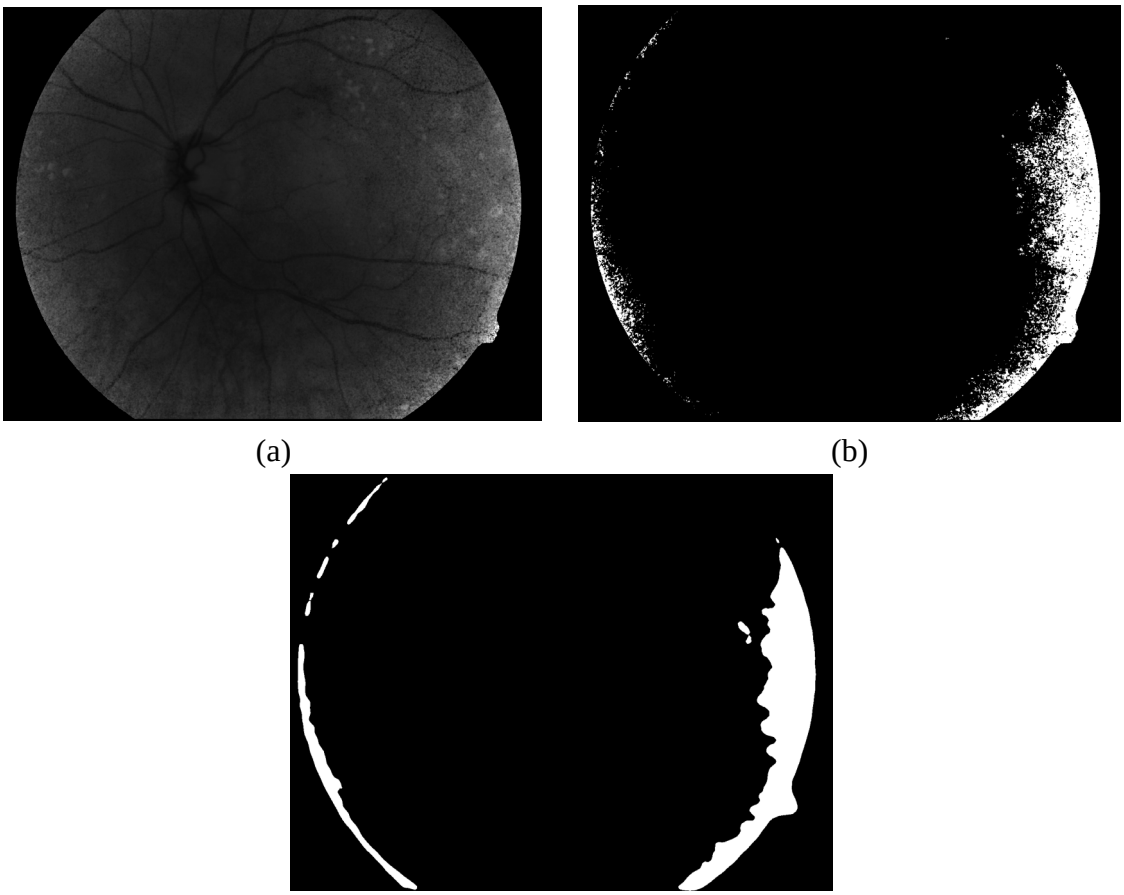

Figure 39: Detecting distorted fundus color in Figure 35: (a) Hue channel (H) in the converted HSV image; (b) Hue channel divided by the intensity channel (V) and thresholded with the preset threshold; (c) Result of using median filtering to (b) showing regions of the fundus having distorted color.

result is morphologically closed in order to smooth the indentations. The result of filling gaps and using morphology operations is shown in Figure 40(b). The figure represents regions of the fundus that either contain too much noise or have distorted color.

Regions of inadequate image quality can be found with the steps described above. A mask containing regions of adequate image quality can be simply created by excluding poor image quality regions from the fundus mask, as shown in Figure 41(a). The illumination-equalized green channel can be masked by the created mask in order to show only regions of the fundus that have adequate image quality. Figure 41(b) shows the result of masking Figure 36(c) by the mask in Figure 41(a). In the result image the illumination is not uneven any more and the image does not contain regions having much noise or fundus pixels whose color is distorted.

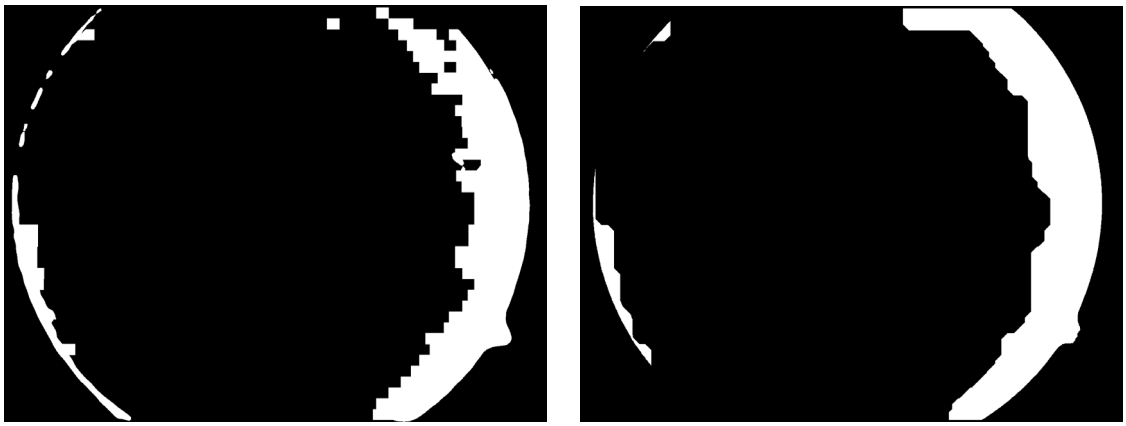

(a) (b)

Figure 40: Combining regions having noise or distorted fundus color: (a) Combination of Figure 38(b) and Figure 39(c); (b) Final mask for regions of poor image quality.

In Figure 42 another example of detecting regions of poor image quality is shown. In this case the color of the fundus is distorted in large regions in the left- and right-hand sides of the fundus as clearly seen in Figure 42(a). The inferiority of the image can also be seen in the green channel of the image, as shown in Figure 42(b), but is more clearly visible when the illumination in the channel is equalized, as shown in Figure 42(c).

Sub-blocks containing more noise pixels than the preset threshold are visible in Figure 42(d). In Figure 42(e) the hue channel (H) has been divided by the intensity channel (V) in the HSV color system and regions where the fundus color is distorted are clearly

(a) (b)

Figure 41: Creating mask of good image quality regions for Figure 35: (a) Binary mask of regions having adequate image quality; (b) Illumination-equalized green channel masked by the binary mask.

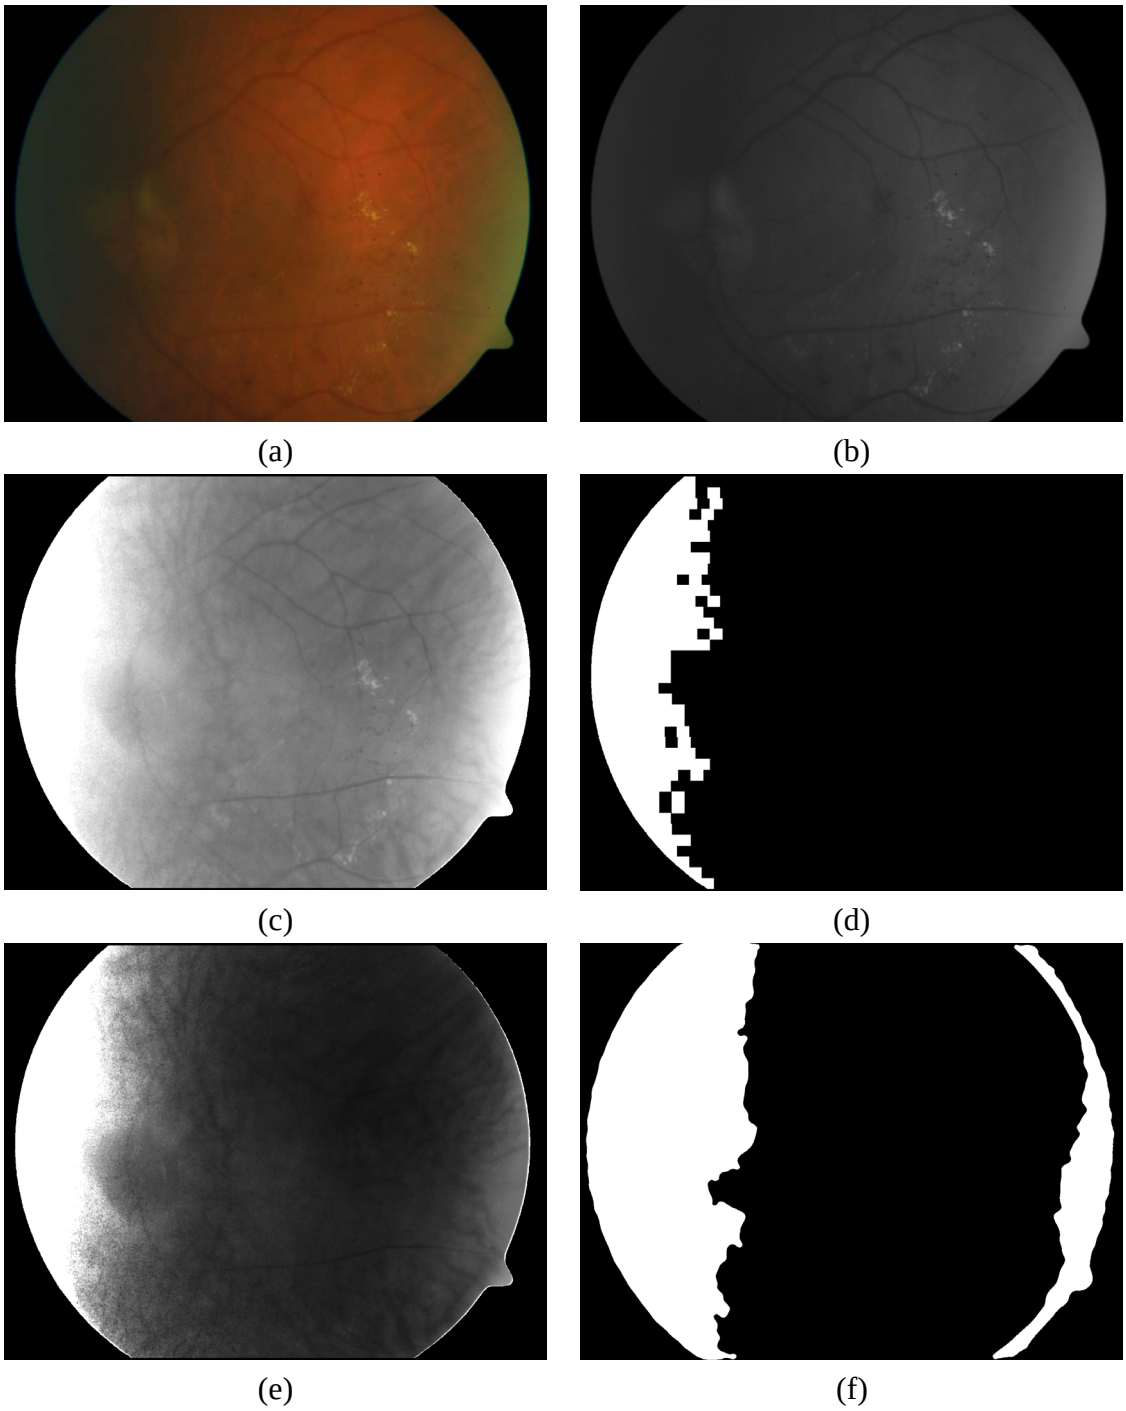

Figure 42: Example of detecting poor image quality regions: (a) Original RGB fundus image; (b) Green channel of the original image; (c) Illumination-equalized green channel; (d) Sub-blocks having more noise pixels than the preset threshold; (e) Hue channel (H) divided by the intensity channel (V) in the HSV color system; (f) Result of thresholding and median filtering (e) representing regions where fundus color is distorted.

seen as bright pixels. Figure 42(f) is the result of thresholding and median filtering Figure 42(e).

The combination of regions having much noise or distorted fundus color is identical to Figure 42(f) since here the regions where the fundus color is distorted cover also regions where there is much noise. Figure 43(a) represents a final mask for adequate image quality. Finally, Figure 43(b) shows the illumination-equalized green channel masked by the adequate image quality mask. Note that the proportion of the fundus image having adequate image quality is relatively low since the fundus color is distorted in an exceptionally large proportion of the fundus.

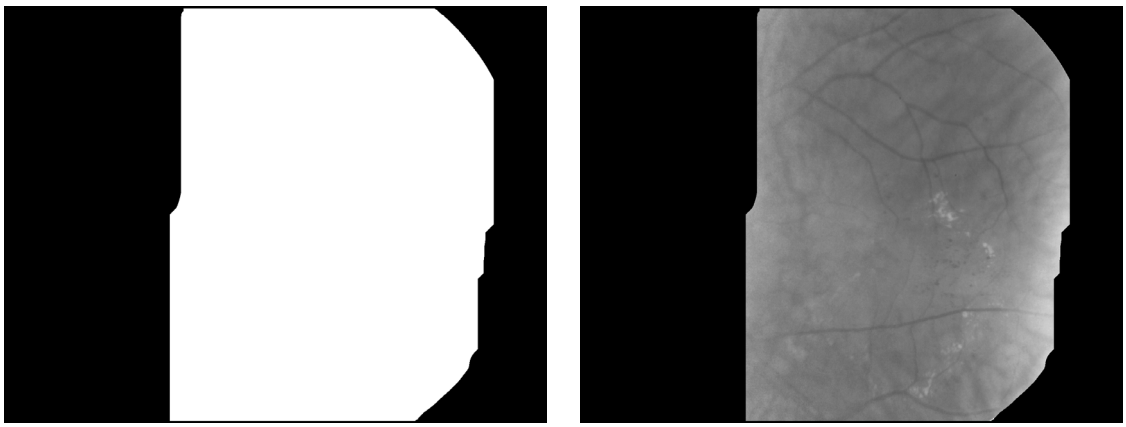

(a) (b)

Figure 43: Result of poor image quality detection: (a) Binary mask for good quality regions; (b) illumination-equalized green channel masked by the binary mask.

Two examples of detecting fundus regions having poor image quality were presented in this sub-section. The final masks specifying regions of adequate image quality, shown in Figure 41(a) and Figure 43(a), have fundus coverages of 85% and 67%, respectively. The mean coverage of adequate image quality masks in the testing image set was 91% in this study.

5 Abnormality detection

The previous section described how the fundus images were acquired and pre-processed in this research. This section explains how abnormal lesions are detected in pre-processed fundus images. In this study the abnormality detection is based on image segmentation using the color information of lesions and on classification of segmented lesions with a rule-based classifier. Image segmentation and object classification techniques were explained in Section 2, and this section describes how these techniques have been used in this study.

The color and brightness of hemorrhages and microaneurysms are relatively similar, as they both consist of blood. The major difference between these two lesion types is the size: microaneurysms are very tiny lesions, whereas hemorrhages are usually much larger. However, as explained above, microaneurysms may be indistinguishable from small hemorrhages, and thus, the term "red small dot" is used for covering both microaneurysms and tiny hemorrhages.

The common feature of hard exudate and soft exudate lesions is that they both appear as brighter than their neighborhood. In this research it was not crucial to find hard exudates and soft exudates separately, it was more important to find all exudates without especially discriminating their type. Thus, hard and soft exudates are considered as one class, exudates. Combining these two evidences is not erroneous, as both should be detected, and the final decision is made by an ophthalmologist.

Exudates appear as brighter and hemorrhages and red small dots dimmer than their neighborhood, not only in the original RGB color fundus images, but also in their green channels. Lesions containing blood (microaneurysms and hemorrhages) are best seen in red-free light [45]. Walter et al. mention that exudates appear most saliently in the green channel of RGB color images [21]. Thus, it was justifiable to concentrate on the illumination-equalized green channel in this research when searching for abnormal lesions.

The detection of all the described abnormal lesion types has been performed in the similar way in this research. The phases of the abnormality detection process are shown in Figure 44. In the first phase the illumination-equalized green channel is segmented into bright regions, dark regions, and background fundus. Candidate hemorrhage and red small dot lesions are searched from dark regions in the second phase. Respectively, candidate exudate lesions are searched from bright regions. Finally in the third phase, candidate lesions

are classified to lesions (hemorrhages, microaneurysms, and exudates) and non-lesions. As a result of the abnormality detection process, lesions of each lesion type are marked in a separate binary mask. Note that abnormality detection is performed for the whole fundus rather than only regions having adequate image quality. This makes it possible to use the abnormality detection results in different ways. For example, an ophthalmologist can select whether all detected lesions are shown or only those that are found in regions having adequate image quality. The next sub-sections describe the three phases of the abnormality detection process more precisely.

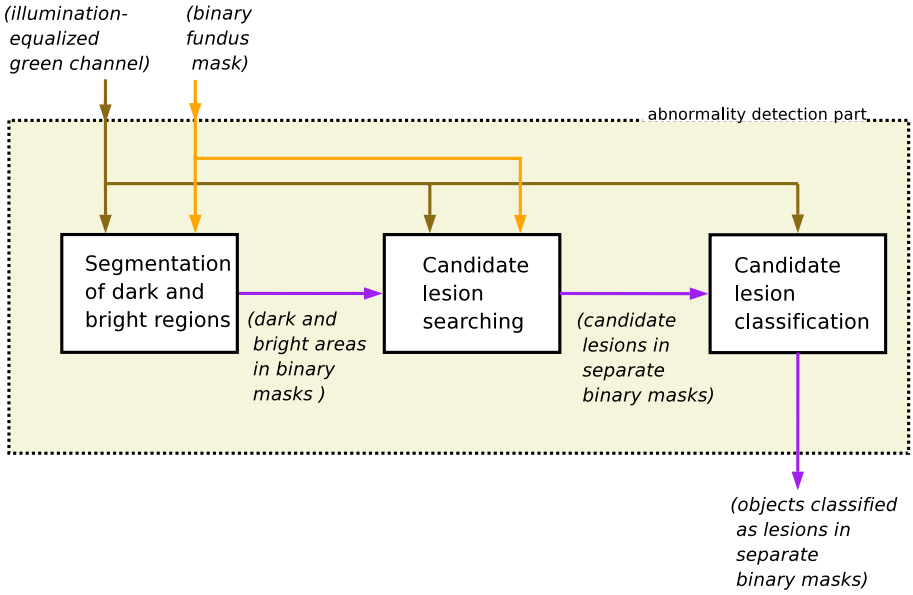

Figure 44: Phases of the abnormality detection process.

5.1 Segmentation of dark and bright regions

The first phase in the abnormality detection process is to find dark regions that contain hemorrhages and microaneurysms and bright regions that contain exudates. Since the color of hemorrhages and microaneurysms is similar to the color of blood vessels, also blood vessels are included in the dark regions. Similarly, the color of the optic disk is the same as that of exudates, and thus the optic disk is included in the bright regions.

Dark and bright regions are found by image segmentation; the green channel where the illumination has been equalized can be segmented into dark regions, bright regions, and background that has medium brightness. Finding dark or bright regions from the illumination-equalized green channel is a trivial task since the proposed illumination

equalization allows the use of global thresholding over the whole image. This is a fundamental result for the success of this work. In illumination-equalized images the brightness of dark or bright lesions or their background is always the same whether an object is near to the center or edge of a fundus.

The threshold values seem to depend linearly on the mean intensity value of the illumination-equalized green channel. This means that a simple equation can be used to to determine the threshold value t for segmenting dark or bright regions as

$$t = km + t_0 \tag{15}$$

where k and t_0 are constants and m is the mean intensity value of the illumination-equalized green channel in the RGB color fundus image.

The constant parameters of the equation, k and t_0, are calculated from the training fundus image set. First a proper threshold is selected for each image in the training set. In this research a simple Matlab-based tool was developed which proposes an initial threshold for an image and allows the user to adjust the threshold. The tool interactively shows the selected threshold and thresholded image, making it possible to find a proper threshold value for the image. When proper thresholds have been selected for all images in the training set, k and t_0 are selected so that the error between the thresholds measured from the images and calculated from Equation 15 is minimized. Error ε between the measured and calculated threshold values can be estimated by using the mean square error criterion

$$\varepsilon = \frac{1}{N}\sum_{i=1}^{N}(T_i - t_i)^2 \tag{16}$$

where N is the number of images, T_i is the measured threshold for the image, and t_i is the calculated threshold for the image. By substituting Equation 15 into Equation 16, Equation 16 can be written as

$$\varepsilon = \frac{1}{N}\sum_{i=1}^{N}(T_i - km_i - t_0)^2 \tag{17}$$

where m_i is the mean intensity of the image. Constants k and t_0 need to be specified only once during the development of the abnormality detection methods. Later a proper threshold for a fundus image can be calculated from Equation 15 simply by substituting the constants and mean intensity value of the image into the equation. Note that actually two sets of constants are needed in the segmentation, one set for specifying the threshold for the dark regions and another for the bright regions.

5.2 Candidate lesion searching

Binary masks containing the dark and bright regions from the segmentation phase and the illumination-equalized green channel are used to search for candidate lesions. Bright regions are selected directly as candidate lesions, since the optic disk is the only normal fundus structure that appears as bright in fundus images. The optic disk is a separate region and can thus be excluded beforehand from the mask of bright regions. However, optic disk detection is not in the scope of this thesis, and thus the optic disk was manually removed from the final result of the detection process.

Searching for candidate lesions from dark regions is more complicated since hemorrhages are not always separate from blood vessels. Thus, it is necessary to remove blood vessels from the dark region mask and after that the remaining regions can be considered as candidate lesions.

It is an easy task to remove the blood vessels if information of the blood vessels is available. There are many studies in the literature concerning blood vessel detection (for example [42] and [46]). Also blood vessel detection was under development in another part of this project, but blood vessel information was not available when this algorithm was developed. Thus, another way for excluding blood vessels had to be used. A morphology-based method was developed for solving the problem.

5.2.1 A morphology-based method for hemorrhage searching

A morphology-based method for hemorrhage detection was developed in this study. The method applies morphology opening and erosion for the binary mask of dark regions produced as the output in the segmentation phase. The opening is performed iteratively with an increasing structuring element until all hemorrhage lesions are separated from blood vessels. The separated objects, or candidate hemorrhage lesions, are marked with 1's and other objects and the background with 0's into a result binary mask.

The morphology-based method is shown in Algorithm 2. The algorithm is processed until all pixels belonging to the original dark regions are marked as candidate hemorrhage lesion pixels or other pixels. In the first step of the algorithm the radius of the disk-shaped structuring element is increased. In the second step the original dark region mask is morphologically opened by the structuring element. In the first iteration round the size of

the structuring element is zero, so no opening is performed. In the next iterations the diameter of the element is always increased by one. The opening is always performed for the original dark region mask. In the third step objects that are too small to be hemorrhage are removed from the opened mask by using morphological erosion. Objects that totally disappear after eroding are considered as too small objects to be hemorrhage, and thus they are moved to a binary mask holding non-hemorrhage objects. The size of the structuring element used in erosion is the same as that used in opening, added by a preset constant (*initialHemorrhageSmallObjectDiskDiameter*). As the size of the structuring element used in removing small objects increases in every iteration round, more and more larger objects are considered as too small to be hemorrhage objects. The rationale for this procedure is that when the structuring element grows, the opening breaks more and more thick blood vessels into smaller parts. If the size of the structuring element used in erosion is not increased during the process, blood vessel parts are considered as proper-sized objects to be candidate hemorrhages. It is important to notice that this algorithm is not used for detecting microaneurysms and small hemorrhages, but small lesion detection is described below. In the next step of the algorithm pixels that have already been marked as candidate hemorrhage lesions or other objects are removed from the mask so that they are not processed again.

Algorithm 2 Morphology-based hemorrhage detection

```
1:  while all dark pixels not marked as candidate hemorrhage or non-hemorrhage do
2:      Increase the size of the structuring element.
3:      Morphologically open the original dark region mask with the structuring element.
4:      Remove small objects from the mask by morphological erosion.
5:      Remove pixels that have already been marked from the mask.
6:      for all separate objects left in the mask do
7:          Calculate properties of the object.
8:          if compactness of the object is higher than a preset threshold then
9:              if the ratio of the object is higher than a preset threshold then
10:                 Mark this object as candidate hemorrhage lesion.
11:             else
12:                 Mark this object as non-hemorrhage object.
13:             end if
14:         end if
15:     end for
16: end while
```

In the fifth step of the algorithm certain properties are calculated for each object in the mask. Hemorrhage lesions have a relatively high compactness value as they are circular-shaped and do not contain holes. If the compactness value for an object is lower than a preset threshold (*minHemorrhageObjectCompactness*), the object is actually a group of

together connected objects and should thus be divided further. If the compactness value for the object is high enough, the ratio of the shortest and longest side of the object is checked. Since hemorrhage lesions are relatively circular, the ratio is near to 1. If the ratio of the object is higher than a preset threshold (*minHemorrhageObjectRatio*), the object is considered to be a candidate hemorrhage lesion, otherwise the object is considered to be a part of blood vessel. Candidate hemorrhage lesions and other objects are marked with 1's in binary masks to exclude already found objects from next iteration rounds. See Subsection 2.3.2 for more information on object compactness and the ratio of dimensions.

When the next iteration round begins, morphological opening is performed for the original dark region mask and the pixels already marked as candidate hemorrhage lesions or other objects are removed from the opened mask. If all pixels have been marked, the processing is stopped.

An example processing for a fundus image containing hemorrhages is shown in Figures 45–48. The original fundus image, its green channel, and illumination-equalized green channel are shown in Figures 45(a–c), respectively. Dark regions containing hemorrhages and blood vessels are obtained by thresholding the illumination-equalized green channel. The result of thresholding is shown in Figure 45(d). The value used for thresholding dark regions is relatively low in order to exclude noisy pixels and background causing the algorithm to fail. A low threshold value usually excludes also small hemorrhages and microaneurysms, and thus a separate red small dot detection algorithm is later used for detecting them.

The mask specifying dark regions is processed iteratively according to the presented algorithm until all objects have been marked as candidate hemorrhage lesions or non-hemorrhage objects. Note that morphological opening is always performed for the original dark region mask presented in Figure 45(d). However, when an object is marked as a candidate hemorrhage lesion or non-hemorrhage object, the object is excluded from the next iteration rounds so that the same objects are not processed multiple times.

Figure 46 shows the result candidate hemorrhage lesions after one, three, and five iterations. Note that all candidate hemorrhage lesions have been found after the fifth iteration. Figure 47 shows pixels that are not yet marked after different iteration rounds. Figure 47(f) shows objects not yet marked after the ninth iteration round, which are finally marked as non-hemorrhages after the tenth round. Then the processing is finished since all the pixels have been marked.

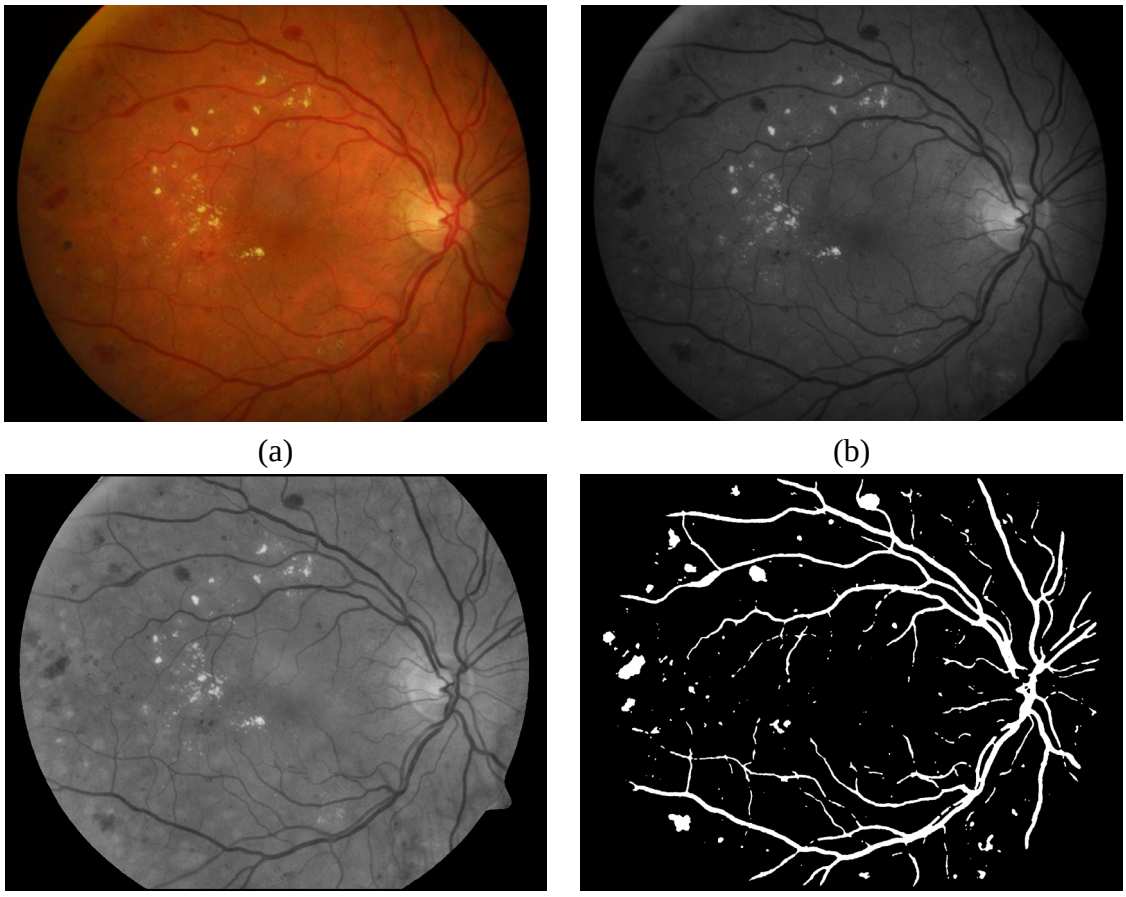

Figure 45: Detecting hemorrhages with the proposed morphology-based method (1/4): (a) Original RGB fundus image; (b) Green channel in the original image; (c) Illumination-equalized green channel; (d) Dark fundus regions obtained by thresholding the illumination-equalized green channel.

The final result of the method is shown in Figure 48, where the found candidate hemorrhage lesions are superimposed on the illumination-equalized green channel. The proposed method is not aimed at small hemorrhage and microaneurysm detection. The next sub-section describes how small lesions are detected from fundus images.

5.2.2 Small lesion searching

The above-proposed hemorrhage detection method cannot be used for detecting small hemorrhages and microaneurysms, and thus a separate method for detecting red small dots is used. The small lesion detection method is also used for improving the sensitivity of exudate searching for small lesion sizes. Both these two lesion types are small and

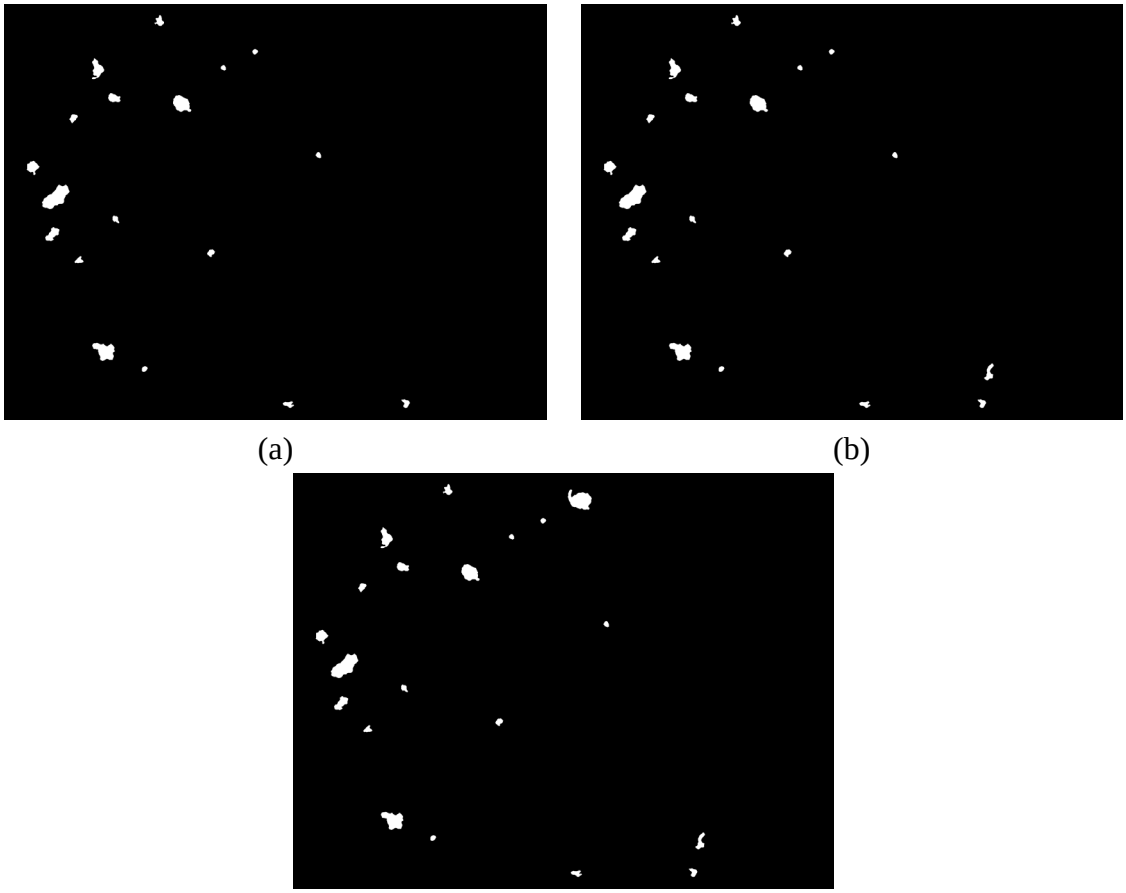

(a) (b)

(c)

Figure 46: Detecting hemorrhages with the proposed morphology-based method (2/4): (a) – (c) Regions marked as candidate hemorrhage lesions after 1, 3, and 5 iterations, respectively.

either darker (red small dots) or brighter (exudates) than their surrounding regions.

The small lesion detection process applies the blob detection method presented in Subsection 2.2.3. The method is performed to the illumination-equalized green channel by using six different filter sizes. The principle of the method is that a small dark object is considered as a candidate red small dot if all its pixels are darker than any of neighboring pixel of the object. Respectively, if all the pixels in the object are brighter than any pixel surrounding the object, the object is considered as a candidate exudate. Two preset parameters were used in blob detection (*minBrightMaskDifference*, *maximumRedMaskDifference*) when specifying the minimum intensity difference between a lesion and its background.

The method is slow since it is performed multiple times with different filter sizes for

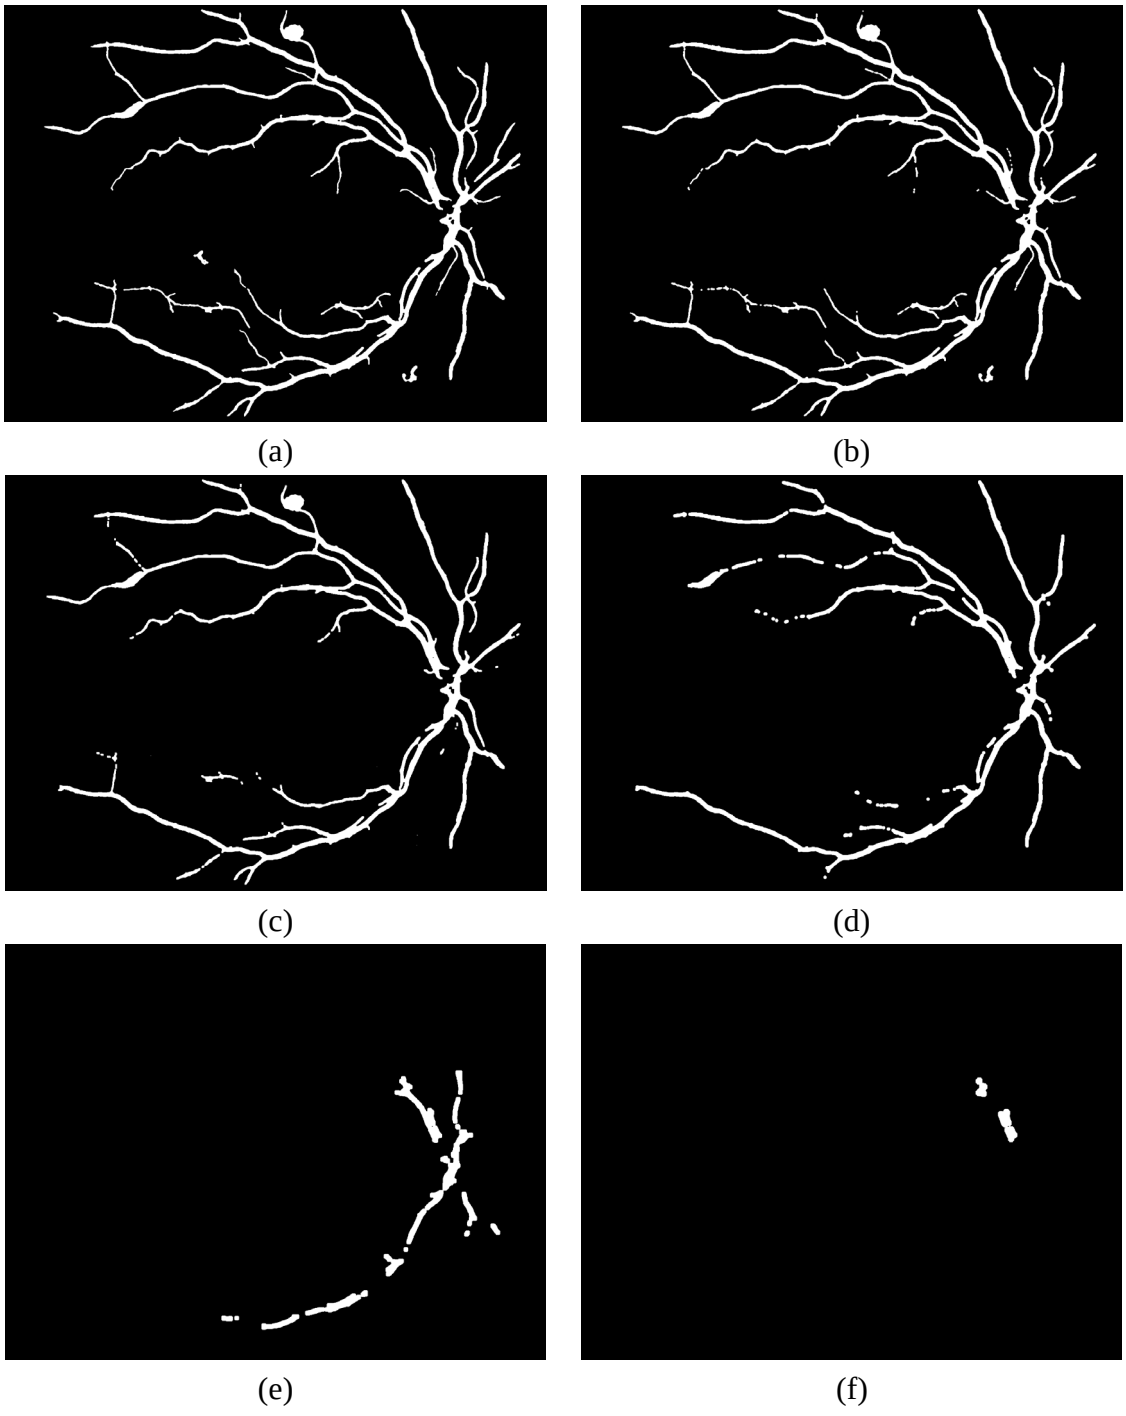

Figure 47: Detecting hemorrhages with the proposed morphology-based method (3/4): (a) – (f) Dark regions not yet marked after 1, 2, 3, 5, 7, and 9 iterations, respectively.

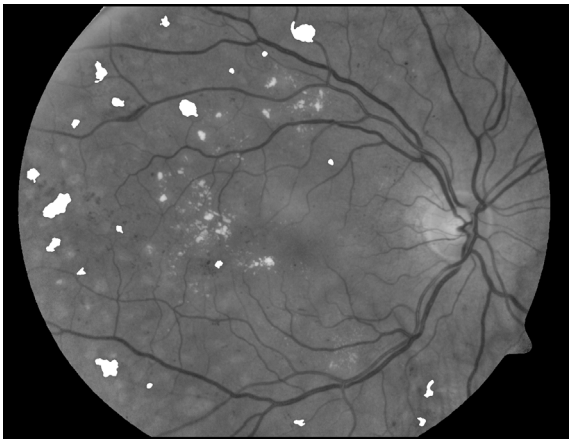

Figure 48: Detecting hemorrhages with the proposed morphology-based method (4/4): Final result of hemorrhage detection superimposed on the illumination-equalized green channel.

every pixel in the image. Thus, small objects were searched only in dark and bright regions produced by image segmentation. Since all small lesions have to be included in the processing, a loose threshold value has to be used when thresholding dark or bright regions. A loose threshold here means a high threshold value when detecting dark regions and a low threshold value when detecting bright regions.

5.3 Candidate lesion classification

After candidate lesions have been found in the previous phase, the illumination-equalized green channel and binary masks containing candidate lesions are given as input to the next phase of the abnormality detection process, candidate lesion classification. In this phase the candidate lesions are classified as lesions or non-lesions. In other words, it is decided whether each candidate lesion is really a lesion or whether it is a normal part of the fundus or noise. The classifier used in this study is rule-based, in other words there is a set of if-else rules that decide whether a candidate lesion is a lesion or a non-lesion. The classifier does not decide the classes of candidate lesions since their classes are already known. In other words, there are separate algorithms for each lesion type in the abnormality detection phase, and thus the classifier only decides whether the candidate lesions are lesions or non-lesions. The classifier uses feature information calculated and measured from candidate lesions as explained in Sub-section 2.3.2.

Since hemorrhages and microaneurysms appear dimmer and exudates brighter than their

background, the most important feature used in the classification is the intensity difference between a candidate lesion and its background in the illumination-equalized green channel. Two parameters (*minExudateMeanRation, minSmallRedAreaMeanRatio*) are used when intensity differences are calculated, as described in Sub-section 2.3.2. See Appendix 1 for a more precise explanation of the parameters.

During classification it is also checked whether the size of a candidate object is within preset limits in pixels. Red small dot and exudate lesions should be larger than preset thresholds (*minRedSmallDotPixelNumber, minExudateObjectPixelNumber*). Candidate hemorrhage lesions from the proposed morphology-based method are classified as hemorrhages if they are larger than a preset threshold (*minPixelsPerHemorrhage*). An extra method was developed to supplement medium-size soft exudate detection. Candidate exudate lesions are classified as exudates if their size is within preset limits (*minPixelsPerSoftExudateObject, maxPixelsPerSoftExudateObject*). The small lesion detection algorithm also finds candidate lesions that are larger than red small dots. Thus, a preset pixel number limit (*maxRedSmallDotPixelNumber*) is used for dividing found small dark lesions into red small dots and hemorrhages.

The rule-based classification is presented in Algorithm 3. The classification is performed for all candidate lesions. First certain properties are calculated and measured from a candidate lesion. If the properties are not between preset limits, the candidate lesion is classified as a non-lesion. If the properties are between preset limits, the final decision is made according to the class of the candidate lesion. If the candidate lesion is a candidate exudate or hemorrhage, it is directly classified as an exudate or hemorrhage lesion, respectively. If the candidate lesion is a candidate red small dot, the classification is done according to the size of the lesion. If the size is smaller than a preset limit, the candidate lesion is classified as a red small dot. If the size is larger than the limit, the candidate lesion is classified as a hemorrhage lesion. The reason for this procedure is that the algorithm that detects small lesions finds, in addition to red small dots, also slightly larger hemorrhages.

It is not in the scope of this research to detect blood vessels and the optic disk in fundus images, but it was done separately in the project. The vessel information is needed when classifying red small dots since the algorithm that searches for red small dots may incorrectly mark some regions of blood vessels as candidate lesions. The information of the optic disk is needed for removing it from the candidate exudate lesions. Algorithm 2 that searches for hemorrhages may sometimes incorrectly mark the fovea of the fundus as hemorrhage since the fovea is a darker circular region in the fundus, as shown

Algorithm 3 Candidate lesion classification

```
1: for all candidate lesions do
2:    Calculate the mean value ratio between the candidate lesion and its neighborhood
      pixels.
3:    Measure the size of the candidate lesion.
4:    if the mean value ratio and the size are not between preset limits then
5:       Classify the candidate lesion as a non-lesion.
6:    else
7:       if lesion type is exudate then
8:          Classify the candidate lesion as an exudate lesion.
9:       else if lesion type is hemorrhage then
10:         Classify the candidate lesion as a hemorrhage.
11:      else if lesion type is red small dot then
12:         if the size of the candidate lesion is smaller than a preset limit then
13:            Classify the candidate lesion as a red small dot.
14:         else
15:            Classify the candidate lesion as a hemorrhage.
16:         end if
17:      end if
18:   end if
19: end for
```

in Figure 49. Thus, the location of the fovea is needed in the study.

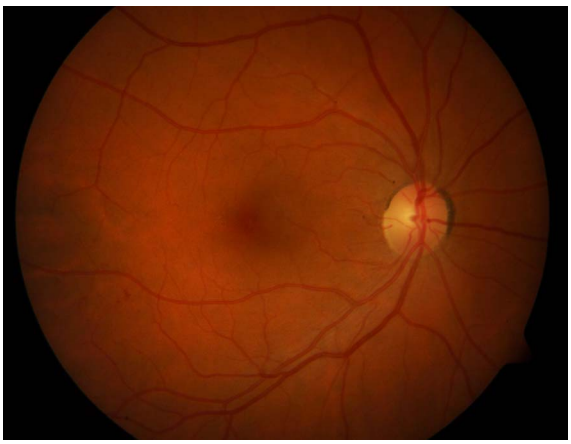

Figure 49: Fovea appearing as a dark region.

After classification, the detected lesions of each lesion type have been marked into a separate binary mask. Binary masks can be further used, for example, for supporting the decision whether some fundus image contains abnormal lesions or for showing the found lesions to the ophthalmologist and saving his/her work. The results can also be used in screening people for diabetic retinopathy.

An example of abnormality detection is presented in Figures 50–52. Figure 50(a) is an original RGB color fundus image. Figure 50(b) represents a created fundus mask where regions of poor image quality have been excluded. Since the illumination has been adequate in the whole image, there are no noisy regions or regions where fundus color is distorted in the image. The green channel of the image and the illumination-equalized green channel masked by the fundus mask are shown in Figures 50(c) and (d) respectively.

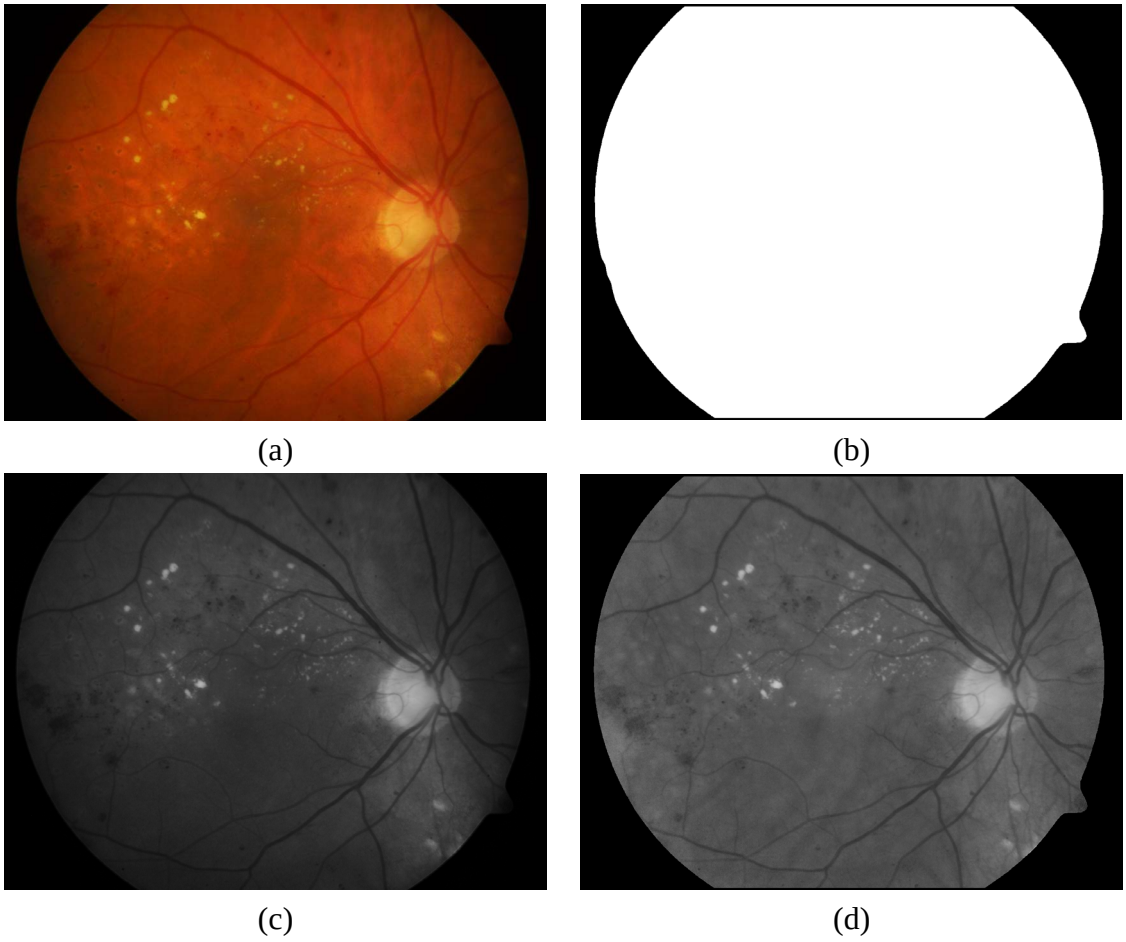

(a) (b)

(c) (d)

Figure 50: Detecting diabetic retinopathy lesions (1/3): (a) Original RGB fundus image; (b) Fundus mask; (c) Green channel of the original image; (d) Illumination-equalized green channel masked by the fundus mask.

The found exudates, hemorrhages, and red small dots are shown in Figures 51(a), (c), and (e), respectively. Note that the red small dots on blood vessels and the optic disk are not excluded since blood vessel and optic disk detection were not studied in this thesis. The found results can be more easily seen if they are superimposed on the illumination-equalized green channel. Figures 51(b), (d), and (f) show the superimposed results of hard exudates, hemorrhages, and red small dots, respectively. Finally all different lesion

types are combined here into the same image for getting the whole picture of the found lesions. Figures 52(a) and (b) show all lesions with different colors superimposed on the illumination-equalized green channel and the original RGB color fundus image, respectively. White color is used for exudates, blue for hemorrhages, and green for red small dots.

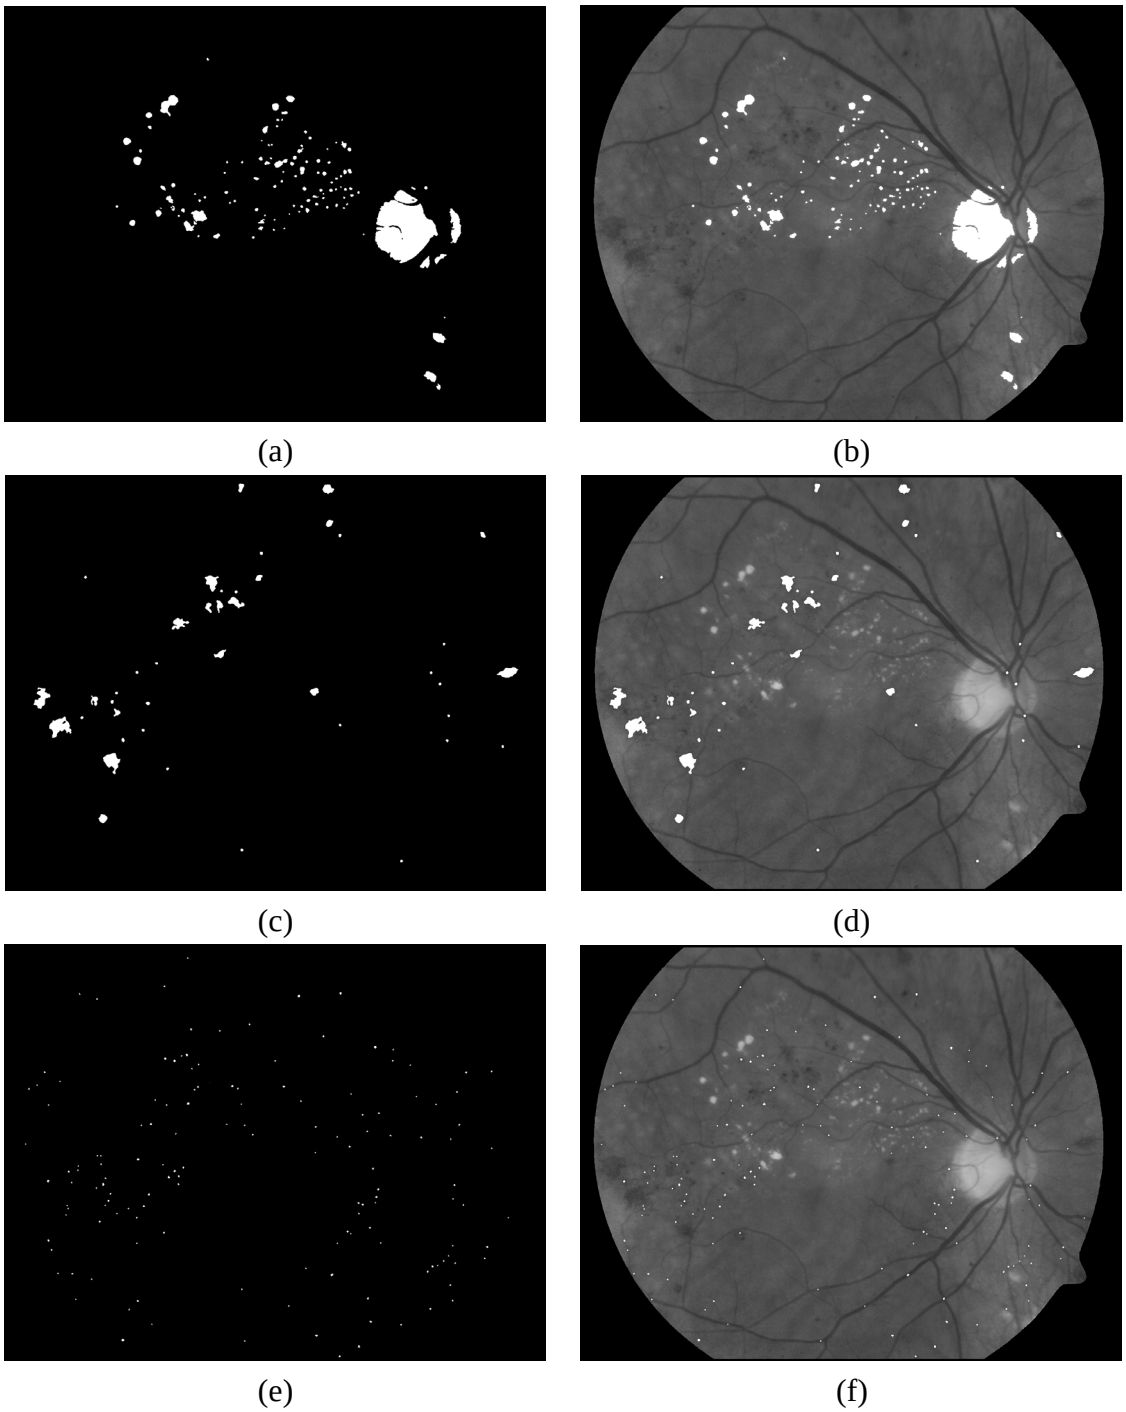

Figure 51: Detecting diabetic retinopathy lesions (2/3), results: (a) Exudates; (b) Exudates superimposed on the illumination-equalized green channel; (c) Hemorrhages; (d) Superimposed hemorrhages; (e) Red small dots; (f) Superimposed red small dots.

6 Experiments and results

6.1 Data and experiments

A total of 173 color fundus images (with a 50-degree field of view) and ground truth information of lesions in the images were provided by the University of Kuopio, Finland. 45 of the images (28 abnormal and 17 normal) were used in algorithm development and a separate image set containing the remaining 128 images (108 abnormal and 20 normal) were used for testing the algorithms. The test set was not used during the development but only once for testing the algorithms and for computing sensitivity and specificity values.

All the 108 abnormal images in the test set contained red small dots. 82 of the images in the set contained hemorrhage. 74 of the images had exudates and 44 did not have them. The remaining 10 images contained funduses where hard exudates were processed by laser photo-coagulation. Since laser scars contain remains of exudates, they are found by exudate detection algorithms. Since patients having laser scars are already diagnosed to have retinopathy, they are not involved in screening. Thus it was considered to be reasonable to exclude the 10 images when calculating the sensitivity and specificity of exudate detection. In addition, 20 fundus images contained neovascularization.

Since the detection of blood vessels and the optic disk were not considered in this re-

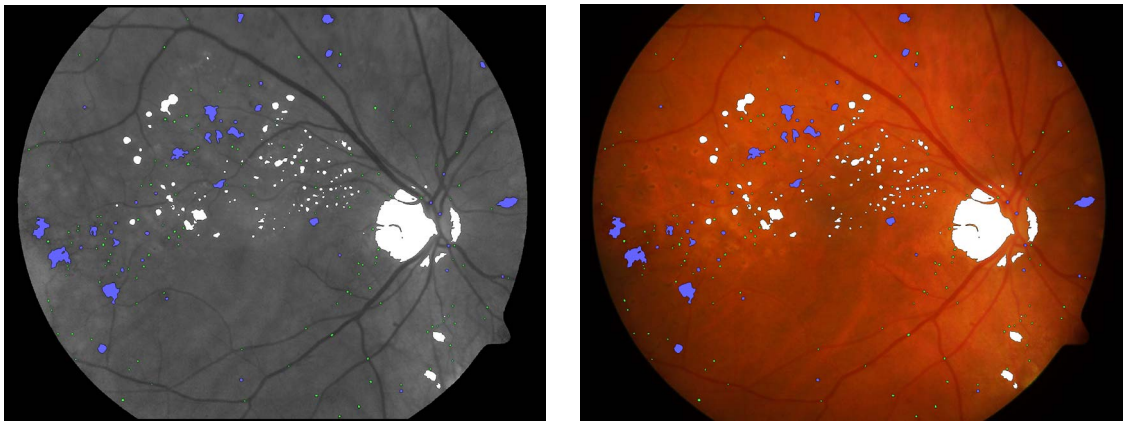

(a) (b)

Figure 52: Detecting diabetic retinopathy lesions (3/3): (a) All detected lesions superimposed on the illumination-equalized green channel; (b) All detected lesions superimposed on the original RGB color fundus image. (white : exudates, blue : hemorrhages, and green : red small dots)

search, but in the other part of the project, the information of blood vessels and optic disk were produced manually. This is reasonable because the purpose of the testing was not to test the whole system (detection of lesions, blood vessels, and optic disk) but to test how well the lesion detection algorithms work in the optimal situation where the blood vessels and optic disks are perfectly detected. The location of the fovea was also given manually for the same reason.

There appear dark small blobs in many images that are not red small dots but dust particles on the camera lens. Dust particles are visible for example in Figures 53(c) and (d). The result images were checked manually and regions of dust particles were excluded, and thus they did not affect the sensitivity and specificity values. Excluding dust particles was considered to be reasonable since in screening the lens should always be clean.

Before the algorithms were tested with the test set, two additional rules were set. They considered the searching for small objects by using a circular filter described in Subsections 2.2.3 and 5.2.2. The detection of red small dots and small exudates is vulnerable to noise, as small lesions may consist of only a few pixels. Since fundus images contain much noise, it may be possible that a few noise pixels are grouped together so that the small lesion detection method incorrectly classifies the group of noise pixels as a lesion. By using a heuristic rule, a fundus image was classified as normal if no more than four red small dots were found in the image. The same rule was used for small exudates with the limit of two. The two rules were applied only to lesions detected by the circular filter method. The rules were not used for other algorithms detecting exudates or other lesion types. It should be noted that the two rules were set according to the training image set.

The testing was performed by running the developed methods in the Matlab 7 environment in Linux Mandrake 10.0 operating system running on an AMD Athlon XP 2600+ processor and 512 MB RAM memory.

6.2 Results and discussion

The developed methods were tested with 128 color fundus images in the test image set. Processing one image took around 6 minutes including pre-processing, image segmentation and object classification. The small object detection took most of the processing time.

The sensitivity and specificity values for the test image set were calculated from Equations 9 and 10. The values were calculated for whole images rather than pixels, which is also common in the literature. The sensitivity and specificity values for each lesion type detection are shown in Table 2.

Table 2: Result sensitivity and specificity percentages for different lesion types.

	Sensitivity %	Specificity %
Red small dots	73	70
Hemorrhages	92	75
Exudates	77	50

The sensitivity and specificity values for the red small dots were 73% and 70%, respectively. Some of the red small dot images marked by the ophthalmologist contained less than five lesions. It was not possible to detect those images since the preset rule specified that an image is classified as normal if the algorithm finds less than five red small dots. The red small dot result is not as good as that reported by Raman et al. [33] (90% sensitivity and 87.5% specificity), but they used only 18 images and the same images were used in training and testing. Also Cree et al. [34] have reported better results (82% sensitivity and 84% specificity), but the results are not comparable with each other since fluorescein angiograms were used in Cree's study, while color fundus images were used in this research.

The sensitivity and specificity values for hemorrhages were 92% and 75%, respectively. The result of hemorrhage detection is better than the only hemorrhage result found in the literature (73.8% sensitivity and 73.8% specificity) [23].

The sensitivity and specificity values for exudates were 77% and 50%, respectively. All the studies found in the literature report better exudate detection results than achieved in this research. The reason for the low sensitivity is that there were relatively many images that contained only a few tiny hard exudate lesions that were many times very difficult to distinguish also with the human eye. An original RGB fundus image with relatively small exudates that were not detected by the developed algorithms is shown in Figure 53(a). It is possible for a person to distinguish exudates from the illumination-equalized green channel shown as magnified in Figure 53(b). However, there were also more difficult images in the testing image set. Figure 53(c) shows an original RGB fundus image with very small exudates. The magnified green channel is shown in Figure 53(d), but it is still difficult to distinguish the exudates.

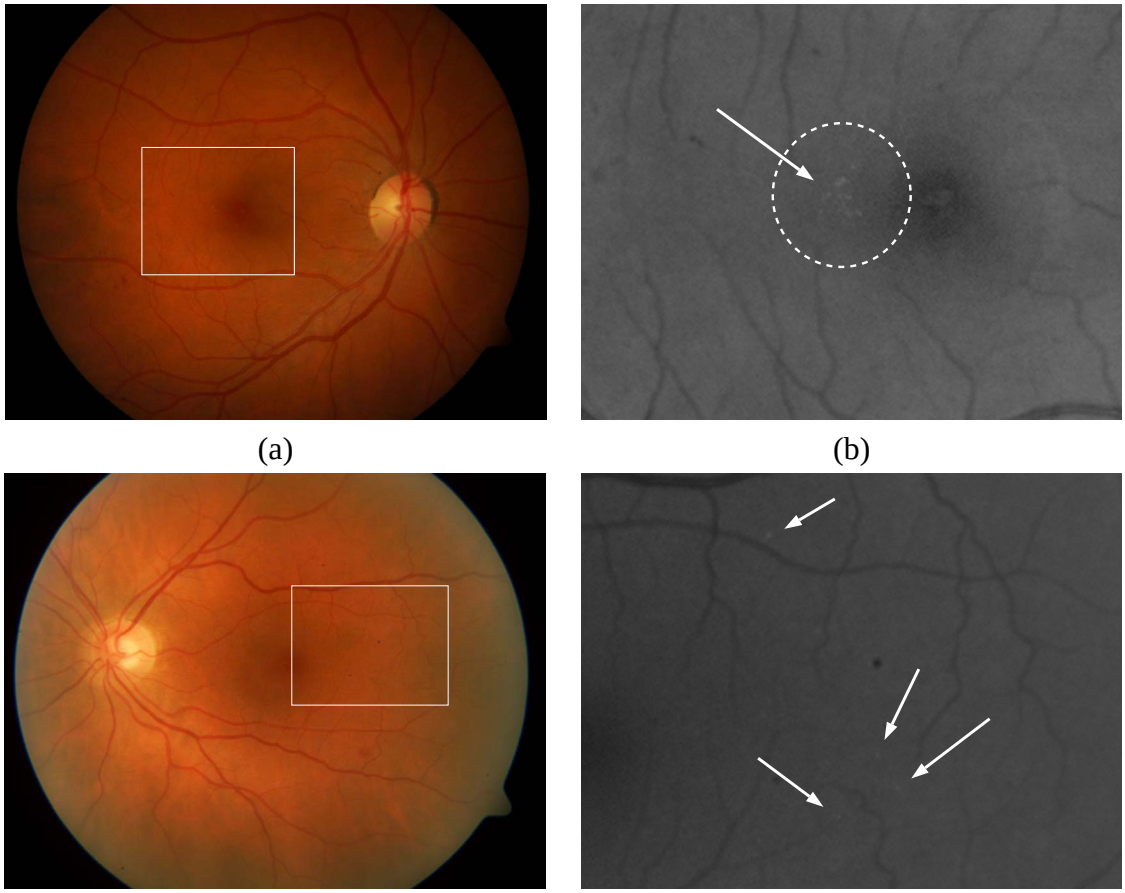

Figure 53: Small exudates not detected by the developed algorithms (a) RGB fundus image containing small exudates; (b) Magnified illumination-equalized green channel; (c) RGB fundus image containing small exudates; (d) Magnified green channel in the original image. The small dark blobs appearing in (c) and (d) are not red small dots but dust particles on the camera lens.

It was also difficult to distinguish soft exudates in several fundus images because the color and intensity of soft exudates were near to the fundus color and intensity. A reason for the low specificity value in exudate detection is that when an image was incorrectly classified to have exudates, there existed only a few small objects in the segmentation results that caused the miss-classification. Also noise appearing in regions that were not excluded by the poor image quality detection methods caused some recognition errors. This was due to the fact that if the illumination equalization method failed in some relatively noisy region, the result region was too bright, causing incorrect image segmentation. However, the accuracy of exudate detection seemed to be good since exudate pixels were accurately classified as exudates and other pixels as non-exudates.

The good exudate detection results found in the literature mean that exudate (especially hard exudate) detection has been studied most. The result of exudate detection reported in this study is not adequate. The algorithms should be better adjusted to achieve sensitivity and specificity values which are good enough. 100% sensitivity and 100% specificity values are unrealistic in the real world, but 90–95% sensitivity and 70–75% specificity would be adequate for most situations, as nearly all fundus images having exudates would be detected, while most of the normal images would be classified as normal.

Due to the lack of resources, neovascularization detection was not studied. However, it was investigated how well images containing neovascularization can be detected by the other developed algorithms. 20 images of the test set contained neovascularization. All those images contained also red small dots, hemorrhages and/or exudates. The developed algorithms detected other lesions in every neovascularization image, and none of the fundus images containing neovascularization was incorrectly classified as a normal image. However, it was not a task of this research to investigate how often other lesion types exist in funduses having neovascularization. If other lesion types usually appear together with neovascularization, the present algorithms would be sufficiently good for not missing neovascularization images. But if neovascularization can also appear alone, a separate algorithm for neovascularization detection should be developed.

7 Conclusions

The goal of this master's thesis was to develop algorithms for detecting different abnormal lesions related to diabetic retinopathy. The lesion types of interest were microaneurysms, hemorrhages, hard exudates, soft exudates, and neovascularization. A method for neovascularization detection was not developed due to the lack of time, however.

A tool for collecting lesion data from the used fundus images was developed by the author and was used by an ophthalmologist who marked lesions in the images. At total of 173 color fundus images and their ground truth were provided by the University of Kuopio, Finland. The images were divided into a training set (45 images) and test set (128 images). The training set was used for developing the algorithms and the test set only for testing the algorithms.

Equalization of uneven illumination was found to be the key issue for the success of the research. Thus, existing illumination equalization methods were compared and the best method was selected. Since abnormal lesions are best visible in the green channel of an RGB color fundus image, the illumination-equalized green channel was used in the abnormality detection process. The abnormality detection process consisted of image segmentation and candidate lesion classification. In addition to thresholding, two novel methods were used in the image segmentation: a circular filter-based method for detecting small lesions and a morphology-based method for hemorrhage detection. Segmented candidate lesions were classified into lesions and non-lesions by using a simple rule-based classifier.

Because the quality of fundus images vary due to difficult imaging conditions, methods for detecting regions of poor image quality were developed. Two types of image quality problems were noticed: noise and distorted fundus color. Both these types appear in regions where illumination has been inadequate. It was not clear why the colors in some images were distorted. It was assumed that all the delivered fundus images were unprocessed. Since the images were taken earlier to be examined by an ophthalmologist, there is a possibility that some of the images were processed to be more suitable for a human examiner. This was supported by the fact that the illumination was somehow uniform in the green channel in those fundus images where the color was distorted. There also existed a couple of red-free fundus images, in other words, RGB fundus images where the red channel values were zero. Images with distorted fundus color prejudiced the result sensitivity and specificity values, as the illumination-equalization methods had problems

with these images.

However, the achieved results were relatively good. Hemorrhages were found almost in every case and also the specificity was good. The microaneurysm detection was moderate and should be improved. The result of the exudate detection was also moderate. There were difficult exudate images in the test set and to cope with them, the exudate detection methods need to be improved. A separate algorithm for neovascularization detection was not developed due to a lack of resources, but all images containing neovascularization were detected by the created algorithms because of the existence of other lesion types. However, the neovascularization detection should be further studied before the algorithms are used in the screening of retinopathy.

The results proved that it is possible to use algorithms for assisting an ophthalmologist to segment fundus images into normal parts and lesions, and thus support the ophthalmologist in his or her decision making. The algorithms detect regions where the image quality is inadequate, and thus it is possible to show to the ophthalmologist what regions are left unprocessed.

It was not investigated how well the developed algorithms can be combined for detecting the existence or absence of retinopathy as a whole. However, it may be possible to achieve better sensitivity and specificity values when searching for retinopathy as a whole and not different lesion types separately, since lesions of different types often appear together. The sensitivity of the combined lesion detection may be higher, because if some of lesion types are missed, there is still a possibility that other lesion types are detected. Also the specificity of the combined system may be high. For example, if only a few small lesions of a certain lesion type, but no other lesion types, are detected, it is very likely that the algorithm yields to a detection error, and thus the image may be classified as normal.

Separate lesion detection algorithms will be combined in the project with algorithms detecting blood vessels, the optic disk, and the location of the fovea. If the result of this combination is sufficiently good, it is possible to automate the screening of retinopathy. In the screening it would be important that the computer-based system has very high sensitivity so that no images having retinopathy are missed. The automated screening system would reduce the work load of ophthalmologists since only fundus images classified as abnormal or unsure by the automated system need to be examined by humans.

There are several proposals for improvements if more resources are available for researching machine vision-based diabetic retinopathy analysis. First, it should be investigated

how the fundus images should be taken. In other words, it should be studied whether it is enough to use a green filter and take gray-scale images containing only the wavelengths of green color or whether a wider range of wavelengths should be used. In this study lesions were searched only in the green channel of RGB color fundus images, but also other color channels may give additional information about lesions. The second proposal concerns image segmentation and classification. The segmentation technique presented in this thesis is not very dynamical, and thus there may appear segmentation errors in some images. It may be necessary to use an image segmentation technique that finds separate object classes from each image and performs the segmentation according to the class information. The third proposal for improvement is that different classifiers (rule-based, statistical, artificial neural networks, etc.) should be compared and the most suitable selected. The fourth proposal is to optimize the used methods to achieve an adequate screening speed. The optimization may also involve the implementation of the developed methods as a stand-alone program, as the current Matlab-based implementation is relatively slow. Finally, the fifth proposal is to use the ground truth of several ophthalmologists instead of a single one. In this study, it was not investigated how much difference there is between diagnoses made by different human screeners, but the ground truth of only one ophthalmologist was used. Since also ophthalmologists may make classification mistakes, it would be worthwhile to use only lesion information that is accepted by several ophthalmologists. When the results of the developed machine vision methods are published, it may be reasonable to mention how much variance there is among human screeners in addition to variance between the computer-based system and a human screener.

Researching diabetic retinopathy detection with machine vision was found to be not only challenging but also rewarding during the process of this thesis. Even if it was not possible to research the topic completely and to select and develop the most optimized methods due to the lack of time, this study gives a ground for further research. In the author's opinion, any work aimed at improving the quality of human life is important.

References

[1] Tekes, national technology agency of Finland. Website (referenced 20th February 2005). http://www.tekes.fi/eng/.

[2] Diabetesliitto: Perustietoa diabeteksestä. Website (referenced 3th February 2005). http://www.diabetes.fi/diabtiet/perus/perus2.htm (In Finnish).

[3] Jack J. Kanski. *Clinical Ophthalmology: A Systematic Approach*. Butterworth-Heinemann, 2003. ISBN 0-7506-5542-9.

[4] *Effective Health Care - Complications of diabetes: Renal disease and promotion of self-management*. Volume 6, number 1. Royal Society of Medicine Press, 2000. ISSN 0965-0288.

[5] Diabetesliitto: Tyypin 1 diabetes. Website (referenced 3th February 2005). http://www.diabetes.fi/diabtiet/perus/tyyppi1.htm (In Finnish).

[6] Diabetesliitto: Tyypin 2 diabetes. Website (referenced 3th February 2005). http://www.diabetes.fi/diabtiet/perus/tyyppi2.htm (In Finnish).

[7] Diabetesliitto: Diabetes Suomessa. Website (referenced 3th February 2005). http://www.diabetes.fi/diabtiet/tiedote/02/tutkijat.htm (In Finnish).

[8] Helsingin ja uudenmaan sairaanhoitopiiri: Yleistä diabetestyypeistä. Website (referenced 3th February 2005). http://www.hus.fi/default.asp?path=1;32;660;546;621;763;1193;5104 (In Finnish).

[9] Annele Eerola, Sirkku Kivisaari, Riikka Eela, and Mikko Rask. *Ikääntyneiden itsenäistä suoriutumista tukeva teknologia*. Eduskunnan kanslian julkaisu, 2001. (In Finnish).

[10] *Effective Health Care - Complications of diabetes: Screening for retinopathy and Management of foot ulcers*. Volume 5, number 4. Royal Society of Medicine Press, 1999. ISSN 0965-0288.

[11] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing, 2nd edition*. Prentice Hall, 2002. ISBN 0-201-18075-8.

[12] John C. Russ. *The Image Processing Handbook, 3rd edition*, pages 196–197. CRC Press LLC and Springer-Verlag GmbH & Co. KG, 1999. ISBN 3-540-64747-3.

[13] Suomen Diabetesliitto ry. *Diabetes ja silmänpohjamuutokset*. Kirjapaino Hermes Oy, 1998. ISBN 951-9406-91-3. (In Finnish).

[14] Suomen Diabetesliiton retinopatiatyöryhmä. Diabeettinen retinopatia: Seuranta, hoito ja näkövammautuneen diabeetikon kuntoutus. *Diabetes-lehden lääkäriliite*, 1992. (In Finnish).

[15] Sight savers: The structure of the human eye. Website (referenced 9th February 2005). http://www.sightsavers.org.uk/html/eyeconditions/human_eye_detailed.htm.

[16] Wikipedia, the free encyclopedia. Website (referenced 9th February 2005). http://en.wikipedia.org/wiki/Main_Page.

[17] Lars B. Bäcklund. *Early Diagnosis of Diabetic Retinopathy by Mass Examinations in Primary Health Care*. PhD thesis, Karolinska Institutet, Sweden, 2000. ISBN 91-628-4224-2.

[18] Daniel E. Singer, David M. Nathan, Howard A. Fogel, and Andrew P. Schachat. Screening for diabetic retinopathy. *Annals of Internal Medicine*, 116:660–671, 1992.

[19] Retinopathy Working Party. A protocol for screening for diabetic retinopathy in europe. *Diabetic Medicine*, 8:263–267, 1991.

[20] D. Usher, M. Dumskyj, M. Himaga, T. H. Williamson, S. Nussey, and J. Boyce. Automated detection of diabetic retinopathy in digital retinal images: a tool for diabetic retinopathy screening. *Diabetes UK. Diabetic Medicine*, 21:84–90, 2003.

[21] Thomas Walter, Jean-Claude Klein, Pascale Massin, and Ali Erginay. A contribution of image processing to the diagnosis of diabetic retinopathy - detection of exudates in color fundus images of the human retina. *IEEE Transactions on Medical Imaging*, 21:1236–1243, October 2002.

[22] Chanjira Sinthanayothin, Viravud Kongbunkiat, Suthee Phoojaruenchanachain, and Apichart Singlavanija. Automated screening system for diabetic retinopathy. In *Proceedings of the 3rd International Symposium on Image and Signal Processing and Analysis*, pages 915–920, 2003.

[23] G. G. Gardner, D. Keating, T.H. Williamson, and A. T. Elliot. Automatic detection of diabetic retinopathy using an articial neural network: a screening tool. *British Journal of Ophthalmology*, 80:940–944, 1996.

[24] Bernhard M. Ege, Ole K. Hejlesen, Ole V. Larsen, Karina Moller, Barry Jennings, David Kerr, and David A. Cavan. Screening for diabetic retinopathy using computer

based image analysis and statistical classification. *Computer Methods and Programs in Biomedicine*, 62:165–175, 2000.

[25] Huan Wang, Wynne Hsu, Kheng Guan Goh, and Mong Li Lee. An effective approach to detect lesions in color retinal images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 181–187, South Carolina, USA, June 2000.

[26] Kheng Guan Goh, Wynne Hsu, Mong Li Lee, and Huan Wang. Adris: an automatic diabetic retinal image screening system. *Medical Data Mining and Knowledge Discovery*, pages 181–210, 2000.

[27] Huiqi Li and Opas Chutatape. Automated feature extraction in color retinal images by a model based approach. *IEEE Transactions on Biomedical Engineering*, 51:246–254, February 2004.

[28] Nicholas P. Ward, Stephen Tomlinson, and Christopher J.Taylor. Image analysis of fundus photographs. *Ophthalmology*, 96:80–86, January 1989.

[29] Russell Phillips, John Forrester, and Peter Sharp. Automated detection and quantification of retinal exudates. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 231:90–94, 1993.

[30] Wynne Hsu, P. M. D. S. Pallawala, Mong Li Lee, and Kah-Guan Au Eong. The role of domain knowledge in the detection of retinal hard exudates. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 246–251, Kauai Marriott, Hawaii, December 2001.

[31] Alireza Osareh. *Automated Identification of Diabetic Retinal Exudates and the Optic Disc*. PhD thesis, Department of Computer Sciene, University of Bristol, 2004.

[32] Alireza Osareh, Majid Mirmehdi, Barry Thomas, and Richard Markham. Comparative exudate classification using support vector machines and neural networks. In *5th International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 413–420, Springer LNCS 2489, September 2002.

[33] B. Raman, E. S. Bursell, M. Wilson, G. Zamora, I. Benche, S. C. Nemeth, and P. Soliz. The effects of spatial resolution on an automated diabetic retinopathy screening system's performace in detection microaneurysms for diabetic retinopathy. In *Proceedings of the 17th IEEE Symposium on Computer-Based Medical Systems (CBMS'04)*, pages 128–133, Bethesda, Maryland, 2004.

[34] M. J. Cree, J. A. Olson, K. C. McHardy, J. V. Forrester, and P. F. Sharp. Automated microaneurysm detection. In *IEEE International Conference on Image Processing*, pages 699–702, Lausanne, Switzerland, 1996.

[35] Mohamed Kamel, Saeid Belkassim, Ana Maria Mendonca, and Aurélio Campilho. A neural network approach for the automatic detection of microaneurysms in retinal angiograms. In *International Joint Conference on Neural Networks*, pages 2695–2699, Washington DC, 2001.

[36] Sherif Abdelazeem. Microaneurysm detection using vessel removal and circular hough transform. In *Nineteenth National Radio Science Conference*, pages 421–426, Alexandria, Egypt, March 2002.

[37] Mohammed Hafez and Sherif Abdel Azeem. Using adaptive edge technique for detecting microaneurysms in fluorescein angiograms of the ocular fundus. In *IEEE Melecon 2002*, pages 479–483, Cairo, Egypt, May 2002.

[38] A. M. Mendonca, A. J. Campilho, and J. M. Nunes. Automatic segmentation of microaneurysms in retinal angiograms of diabetic patients. In *10th International conference on Image analysis and processing*, pages 728–733, Venice, Italy, September 1999.

[39] Zheng Liu, Chutatape Opas, and Shankar M. Krishnan. Automatic image analysis of fundus photograph. In *Proceedings - 19th International Conference - IEEE/EMBS*, pages 524–525, Chicago, IL, USA, October 1997.

[40] Gang Luo, Opas Chutatape, Huiqi Li, and Shankar M. Krishnan. Abnormality detection in automated mass screening system of diabetic retinopathy. In *Fourteenth IEEE Symposium on Computer-Based Medical Systems*, pages 132–137, Bethesda, Maryland, 2001.

[41] Huiqi Li and Opas Chutatape. A model-based approach for automated feature extraction in fundus images. In *Ninth IEEE International Conference on Computer Vision Volume 1*, pages 394–399, Nice, France, 2003.

[42] Chanjira Sinthanayothin, James F. Boyce, Helen L. Cook, and Thomas H. Williamson. Automated localisation of the optic disc, fovea, and retinal blood vessels from digital colour fundus images. *Br J Ophthalmol*, 83:902–910, 1999.

[43] Alireza Osareh, Majid Mirmehdi, Barry Thomas, and Richard Markham. Automatic recognition of exudative maculopathy using fuzzy C-means clustering and neural

networks. In *Medical Image Understanding and Analysis*, pages 49–52, BMVA Press, July 2001.

[44] Alireza Osareh, Majid Mirmehdi, Barry Thomas, and Richard Markham. Classification and localisation of diabetic-related eye disease. In *7th European Conference on Computer Vision*, pages 502–516, Springer LNCS 2353, May 2002.

[45] Timo Hellstedt, Eija Vesti, and Ilkka Immonen. Identification of individual microaneurysms: a comparison between fluorescein angiograms and red-free and colour photographs. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 234:13–17, 1996.

[46] Joes Staal, Michael D. Abramoff, Meindert Niemeijer, Max A. Viergever, and Bram van Ginnesen. Ridge-based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging*, 23:501–509, April 2004.

Appendix 1. Parameters of the developed algorithms

The parameters used in the developed algorithms are listed in the table below. Note that all used gray-level parameters are within range $[0, 1]$ or $[-1, 0]$.

Parameter name	Default value
clearEdgeThickness	6
badHueChannelThreshold	0.3
randomNoiseBlockWidth	35
randomNoiseBlockHeight	30
randomNoiseSensitivity	0.025
randomNoiseLimit	0.1
minExudateObjectPixelNumber	10
minExudateMeanRatio	1.2
minimumBrightMaskDifference	3/255
minPixelsPerSoftExudateObject	200
maxPixelsPerSoftExudateObject	2000
initialHemorrhageSmallObjectDiskDiameter	6
minHemorrhageObjectCompactness	0.5
minHemorrhageObjectRatio	1/2.5
minPixelsPerHemorrhage	150
maximumRedMaskDifference	-2/255
minSmallRedAreaMeanRatio	1.15
minRedSmallDotPixelNumber	10
maxRedSmallDotPixelNumber	40

The first block of the parameters is used for detecting and removing noise as follows

clearEdgeThickness specifies how many pixels are removed from the fundus edge in the fundus mask. The removing is done because the edge of fundus is very noisy in the fundus images.

badHueChannelThreshold The hue channel (H) in the HSV color system is divided by the intensity channel (V) and the result is thresholded by this value for finding regions of the image where the colors are spurious. The lower the value, the more sensitively the regions having bad colors are excluded from the fundus mask.

randomNoiseBlockWidth A fundus image is divided into sub-blocks when searching for regions having too much noise. This value is the initial width of one sub-block.

(to be continued)

Appendix 1. (continues)

randomNoiseBlockHeight Same as the previous one, but now the value means the initial height of one sub-block.

randomNoiseSensitivity is used for finding pixels that are considered as noise. If the difference between an original image pixel and the median filtered pixel is higher than this value, the pixel is considered to be noise. The lower the value, the more sensitively a pixel is classified as a noise pixel.

randomNoiseLimit is used for deciding whether a sub-block contains enough noise to be classified as a noise block. For example, the value 0.1 means that if more than 10% of the pixels in a sub-block are noise, the whole sub-block is marked to be noise.

The second block of parameters is used for classifying candidate exudate lesions into exudate lesions and non-exudate lesions as follows

minExudateObjectPixelNumber is the minimum number of pixels that are considered as an exudate lesion.

minExudateMeanRatio is the minimum ratio in the mean values of an exudate lesion and its neighboring pixels. For example, 1.2 means that the mean value of an exudate lesion divided by the mean value of pixels around the lesion should be at least 1.2.

minimumBrightMaskDifference A small candidate exudate lesion is classified as an exudate lesion if the darkest pixel of the lesion is brighter than the brightest pixel around the lesion. This value tells how much brighter the darkest lesion pixel should be when compared to the brightest neighboring pixel.

minPixelsPerSoftExudateObject is the minimum number of pixels that are considered as a soft exudate lesion when soft exudates are searched with an extra algorithm. The extra algorithm finds medium size exudates that appear as relatively dim lesions.

maxPixelsPerSoftExudateObject The same as the previous one, but now the maximum size of soft exudate lesions searched by the extra algorithm is specified.

(to be continued)

Appendix 1. (continues)

The third block of parameters is used for finding candidate hemorrhage lesions and classifying them into hemorrhage lesions and non-hemorrhage lesions as follows

initialHemorrhageSmallObjectDiskDiameter [1] is the starting diameter of a disk-shaped structuring element used for removing small objects that are out of interest when searching for hemorrhages by a morphology-based method. Note that the parameter is not actually a real diameter but a variable used in structuring elements in the Matlab program.

minHemorrhageObjectCompactness Hemorrhage lesions are compact, which means that there are no holes inside the hemorrhage lesions. A dark object is not divided into smaller parts any more if its compact value is higher than this parameter. The compact value is calculated so that the number of lesion pixels is divided by the number of pixels of such a rectangle where the lesion fits precisely.

minHemorrhageObjectRatio Hemorrhage appears as round or oval lesions, but not as elongated objects. The ratio is calculated so that the shortest axis of the object is divided by the longest axis. If the ratio of a candidate lesion is higher than the parameter, the candidate lesion is classified as hemorrhage.

minPixelsPerHemorrhage is the minimum number of pixels in a hemorrhage lesion. Note that the algorithm using this parameter is not used for detecting small hemorrhages.

(to be continued)

[1] The name of this parameter was earlier "initialHemorrhageOpeningDiskDiameter" but was changed because the previous name was specious.

Appendix 1. (continues)

The last block of parameters is used for detecting red small dots and those small hemorrhages that are not searched in the previous block as follows

maximumRedMaskDifference A small candidate red lesion will be classified as a lesion if the brightest pixel of the lesion is darker than the darkest pixel around the lesion. This value tells how much darker the brightest lesion pixel should be when compared to the darkest neighboring pixel.

minSmallRedAreaMeanRatio is the minimum ratio in the mean values of the neighboring pixels of a red lesion and the lesion pixels. For example, 1.15 means that the mean value of the pixels around the lesion divided by the mean value of the lesion pixels should be at least 1.15.

minRedSmallDotPixelNumber is the minimum number of pixels in a red small dot lesion. If a candidate lesion has less pixels, it is considered as noise.

maxRedSmallDotPixelNumber is the limit between red small dots and hemorrhages. If the lesion has less pixels than the set parameter, the lesion is considered as a red small dot. Otherwise the lesion is considered as a hemorrhage.