



HAL
open science

Image processing methods for computer-aided screening of diabetic retinopathy

Xiwei Zhang

► **To cite this version:**

Xiwei Zhang. Image processing methods for computer-aided screening of diabetic retinopathy. Other. Ecole Nationale Supérieure des Mines de Paris, 2014. English. NNT : 2014ENMP0024 . tel-01083819

HAL Id: tel-01083819

<https://pastel.hal.science/tel-01083819v1>

Submitted on 18 Nov 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

École doctorale n°432 :
Sciences des Métiers de l'Ingénieur

Doctorat ParisTech

T H È S E

pour obtenir le grade de docteur délivré par

l'École nationale supérieure des mines de Paris

Spécialité « Morphologie Mathématique »

présentée et soutenue publiquement par

Xiwei ZHANG

le 4 juillet 2014

**Image processing methods for computer-aided screening of
diabetic retinopathy**

Directeur de thèse : **Etienne DECENCIÈRE**

Jury

M. Nicolas PASSAT , Professeur, CReSTIC, Université de Reims Champagne-Ardenne	Rapporteur
M. Hugues TALBOT , Professeur, A2SI, ESIEE	Rapporteur
M. Guy CAZUGUEL , Docteur, LaTIM, Télécom Bretagne	Examinateur
M. Bruno LAÏ , Docteur, ADCIS	Examinateur
M. Jesús ANGULO , Maître de recherche, CMM, MINES ParisTech	Examinateur
M. Etienne DECENCIÈRE , Maître de recherche, CMM, MINES ParisTech	Directeur de thèse

MINES ParisTech

Centre de Morphologie Mathématique, Mathématiques et Systèmes

35, rue Saint-Honoré, 77305 Fontainebleau

**T
H
È
S
E**

Acknowledgments

I would like to thank my family and friends for their supports and encouragements. I would like to thank my advisor, Prof. Dr. Jean-François Roy, for his guidance and support during my PhD journey. I would like to thank my colleagues and friends for their help and support during my PhD journey.

First of all, I would like to thank Nicolas Passat, Hugues Talbot, Guy Cazuguel, Bruno Laÿ and Jesús Angulo for precious time and advice, and for being members of my dissertation committee.

I would like to thank Etienne Decencière, the supervisor of my thesis. It has been a great pleasure and honor to work with him. He created a harmonious advisor-advisee relation and a well-organized working environment, which are the key points for a successful thesis. He provided me with an autonomous working environment, while giving me the important advice and inspirations. I would like to thank Guillaume Thibault, who worked as post-doc in our project, for giving me precious advice on my research and methodology with great enthusiasm.

I would like to thank Fernand Meyer, the former director of the Centre for Mathematical Morphology (CMM), for the advices on my work and for the time spent on reviewing my thesis. I would like to thank Catherine Moysan, the secretary of CMM, for providing help in every aspect of my work. I would like to thank Vincent Morard, former PhD student in the center, for his advice on my work. I would like to thank Jean-Claude Klein, co-advisor during my internship in CMM before my thesis. I would like to extend my gratefulness to all members of CMM, Michel Bilodeau, Dominique Jeulin, Serge Beucher, Jesús Angulo, Beatriz Marcotegui, Petr Dokladal, Matthieu Faessel, Francois Willot and Serge Koudoro, and all other PhD students and interns.

I would like to thank all the partners of the TELEOPHTA project, Ali Erginay, Pascale Massin, Agnès Chabouis and Zeynep Victor from APHP-Hôpital Lariboisière, Guy Cazugeul, Gwénoél Quéllec and Mathieu Lamard from Télécom Bretagne and LaTIM, Bruno Laÿ, Ronan Danno and Damien Elie from ADCIS. It has been a great time to work with them.

I would like to extend my gratefulness to Thomas Walter. Even if I met him only towards the end of my PhD thesis, his thesis gave me lots of great inspirations on my work.

In the end, I would like to thank my parents in China and my friends for their supports. They make me stronger and help me going further in my life.

During my thesis, not only have I learnt specific scientific knowledge in my research domain, but also the methodologies of research and even the way to think. Thanks again to everyone who contributed to this work.

Diabetic retinopathy is the main cause of blindness among the middle-aged population. An early detection and adapted treatment considerably reduce the risk of sight loss. Medical authorities recommend an annual examination to diabetic patients. Several diabetic retinopathy screening programs have been deployed to enforce this recommendation. The aim of the TeleOphta project was to automatically detect normal examinations in a diabetic screening system, in order to reduce the burden on readers, and therefore serve more patients. This thesis proposes several methods to extract information linked to diabetic retinopathy lesions from color eye fundus images.

The detection of exudates, microaneurysms and hemorrhages is discussed in detail. One of the main challenges of this work is to deal with clinical images, acquired by different types of eye fundus cameras, by different persons. Therefore the data base heterogeneity is high. New pre-processing methods, which perform not only normalization and denoising tasks, but also detect reflections and artifacts in the images, are proposed. Novel candidate segmentation methods based on mathematical morphology, and new textural and contextual features for lesion characterization, are proposed. A random forest algorithm is used to detect lesions among the candidates. The proposed methods make extensive use of new residue analysis methods.

Moreover, three new publicly available retinal image databases, e-ophta EX, e-ophta MA and e-ophta HM, respectively designed to develop and evaluate exudate, microaneurysms and hemorrhages detections methods, are proposed in this work. The images are extracted from the OPHDIAT telemedicine network for diabetic retinopathy screening. Manual annotations of the lesions are given in detail in these databases. The proposed algorithms are evaluated on these databases.

The proposed methods have been integrated within the TeleOphta system, which is presented and evaluated on two large databases. Each patient record is classified into two categories: “To be referred” or “Not to be referred”. The classification is based not only on the results of the presented methods, but also on image signatures provided by other partners, as well as on medical and acquisition-related information. The evaluation shows that the TeleOphta system can make about 2 times more patients benefit from the diagnosis service, based on the existing telemedicine network.

Keywords: Image processing, mathematical morphology, pattern recognition, diabetic retinopathy screening.

Abstract

La rétinopathie diabétique est la cause principale de cécité dans la population en âge de travailler. Une détection précoce et un traitement adapté permettent de réduire considérablement le risque de perte de vue. Les autorités médicales recommandent un examen annuel pour les patients diabétiques. Plusieurs programmes de dépistage de la rétinopathie diabétique ont été déployés pour appliquer cette recommandation. L'objectif du projet TeleOphta était de détecter automatiquement des examens normaux dans un système de dépistage du diabète, afin de réduire le fardeau des lecteurs, et donc servir plus de patients. Cette thèse propose plusieurs méthodes pour extraire des informations liées à des lésions provoquées par la rétinopathie diabétique dans des images en couleurs du fond d'œil.

La détection des exsudats, microanévrismes et hémorragies est discutée en détail. L'un des principaux défis de ce travail est de traiter des images cliniques, acquises avec différents types de caméras de fond d'œil, par des personnes différentes. Par conséquent, l'hétérogénéité de la base de données est élevé. Des nouvelles méthodes de pré-traitement, qui effectuent non seulement des tâches de normalisation et de débruitage, mais aussi de détection de réflexions et d'artefacts optiques, sont proposées. Des méthodes de segmentation des candidats basées sur la morphologie mathématique, et de nouveaux descripteurs de texture et de contexte sont proposées pour la caractérisation des lésions. Un algorithme de forêts aléatoires est utilisé pour choisir les lésions parmi les candidats. Les méthodes proposées utilisent largement des nouvelles méthodes d'analyse des résidus.

En outre, trois nouvelles bases de données publiques d'images de la rétine, e-ophta EX, e-ophta MA et e-ophta HM, respectivement conçues pour développer et évaluer les méthodes de détection d'exsudats, de microanévrismes et d'hémorragies, sont proposées dans ce travail. Les images ont été extraites du réseau de télémédecine OPHDIAT pour le dépistage de la rétinopathie diabétique. Des annotations manuelles détaillées des lésions sont fournies avec ces bases de données. Les algorithmes proposés sont évalués sur ces bases.

Les méthodes proposées ont été intégrées dans le système TeleOphta, qui est présentée et évaluée sur deux grandes bases de données. Chaque dossier du patient est classé en deux catégories: "Pour avis" ou "Normal". La classification est basée non seulement sur les résultats des méthodes présentées, mais aussi sur les signatures d'image fournies par d'autres partenaires, ainsi que sur l'information médicale du patient, et les données liées à l'acquisition. L'évaluation montre que le système TeleOphta permet de traiter deux fois plus de patients dans un réseau de dépistage, à moyens constants.

Mots clés: Traitement d'images, morphologie mathématique, reconnaissance de formes, dépistage de la rétinopathie diabétique.

1	Introduction	9
1.1	TeleOphta Project	9
1.2	Medical background	11
1.3	e-ophta database	14
1.4	Research proposal	15
2	Residue analysis methods	17
2.1	Pre-processing and residue operator	17
2.2	Multi-threshold decomposition	23
2.3	Conclusions and perspectives	27
3	Spatial calibration of fundus image	29
3.1	Possible approaches	30
3.2	Spatial calibration using the width of the field-of-view	30
3.3	Validation	33
3.4	Conclusions and perspectives	33
4	Exudates detection	35
4.1	Introduction	35
4.2	State of the art	37
4.3	Pre-processing	38
4.4	Candidates extraction	42
4.5	Classification and risk evaluation	43
4.6	Results	51
4.7	Conclusion	57
5	Microaneurysms detection	59
5.1	Introduction	59
5.2	Microaneurysm characteristics	59
5.3	Microaneurysms detection method	62
5.4	Method evaluation	75

CONTENTS

5.5	Conclusions and perspectives	79
6	Hemorrhages detection	81
6.1	Introduction	81
6.2	Main difficulties	81
6.3	Hemorrhages detection method	83
6.4	Method evaluation	95
6.5	Conclusions and perspectives	97
7	TeleOphta system evaluation	101
7.1	Heterogeneous information fusion	101
7.2	Results	102
7.3	Conclusions	104
8	Conclusions and perspectives	107
8.1	Contributions	107
8.2	Perspectives	109
A	Field of view segmentation	111
B	Vessel segmentation	117
B.1	State of the art	117
B.2	Method based on alternating sequential filter	119
B.3	Method based on Gabor filter	120
B.4	Segmentation results on bad quality images	124
C	Optic Disc localization	129
C.1	State of the art	129
C.2	Border reflection detection	130
C.3	Intensity and size information	131
C.4	Vessel analysis	132
C.5	Selection	134
C.6	Evaluation and conclusions	134
	Glossary	137
	References	140

1.1	OPHDIAT telemedicine network structure and TeleOphta project.	10
1.2	TeleOphta strategy. Workflow of Patient record classification.	11
1.3	Anatomic image of the eye and fundus image.	12
1.4	Different capture fields of fundus images.	12
1.5	Clinical signs of diabetic retinopathy	14
1.6	Fundus images examples from the e-ophta database.	15
1.7	Screenshot of TeleOphta annotation software.	16
2.1	Pre-processing algorithms.	19
2.2	Residues operators.	20
2.3	Residue operators applied on other images.	22
2.4	Thresholding comparison of a pre-processed image	23
2.5	Thresholding of a microaneurysm.	24
2.6	Multi-threshold on an exudate.	25
2.7	Max-tree decomposition and node selection.	25
2.8	Layer width as a function of the Max-tree layer.	26
3.1	Images of different resolutions.	29
3.2	Normalization based on the distance between the OD and the Fovea.	31
3.3	Capture angle of fundus image.	31
3.4	Image contains a dark zone and the segmented FOV.	32
3.5	Spatial normalization by estimations of the sizes of retinal structures and lesions.	32
3.6	Spatial calibration applied on images with different capture angles.	34
4.1	Color retinal image and manual annotation	35
4.2	Examples of fundus images in e-ophta EX.	36
4.3	Dark structures removal based on morphological inpainting.	39
4.4	Three kinds of bright structures.	40
4.5	Adaptive template for the removal of bright structures.	41
4.6	Reflection level estimation and computation of preprocessed image.	42
4.7	Bright border regions segmentation.	43

LIST OF FIGURES

4.8	Large exudate candidates extraction illustration.	44
4.9	Precise shape of candidates extraction and small candidates.	45
4.10	Optical artifacts.	46
4.11	Illustration of computing the swamping feature.	47
4.12	Illustration of computing the swamping feature on fundus image.	48
4.13	Local variance and ultimate opening	49
4.14	Reflections beside vessels	49
4.15	Contextual features computed on neighbor structures	50
4.16	Features importance (exudates)	51
4.17	Risk function.	52
4.18	Definition of True Positive and False Positive pixels.	53
4.19	Sensitivity and precision for pixel level evaluation.	55
4.20	Example of pixel level validation.	55
4.21	ROC curves and AUC values of exudates detection on e-optha EX.	56
4.22	ROC curves and AUC values of exudates detection on other databases.	57
5.1	A microaneurysm in color eye fundus image and its green channel.	60
5.2	Reconstruction of MA.	61
5.3	Details of a microaneurysm.	62
5.4	Contrast distribution of microaneurysms.	63
5.5	Microaneurysm size distribution.	64
5.6	Microaneurysm models.	65
5.7	A dark region among exudates leads to a rise in the pre-processed image.	66
5.8	Maxima over a high platform removed by a reconstruction.	66
5.9	Detection of microaneurysms (I).	67
5.10	A microaneurysm over a vessel.	68
5.11	Detection of microaneurysms (II).	69
5.12	Disconnected vessels become false positives.	71
5.13	Disconnected vessel fragments and microaneurysm near a vessel comparison.	72
5.14	Schematic illustration of vessel fragments and true microaneurysm.	73
5.15	Selection of region-of-interest by vessel orientation map.	73
5.16	Features importance (microaneurysms).	74
5.17	Classification result (microaneurysms).	75
5.18	Classification results superposed with expert's annotation (microaneurysms).	76
5.19	FROC curves for different sets of features (microaneurysms).	77
5.20	ROC curve calculated on the entire e-optha MA database (using TOP-20 features).	78
6.1	Examples of hemorrhages.	82
6.2	Main sources of mis-detected hemorrhages.	83
6.3	Hemorrhages in the macula region.	83
6.4	Flow chart of the hemorrhages candidates detection method.	85
6.5	Bright structures removal by morphological inpainting.	86
6.6	Influence of border reflections on background removal.	86
6.7	Thresholding to get first mask.	87
6.8	Reconstruction by raising global intensity.	88
6.9	Vessel circular variance analysis.	89

6.10	Details of vessel orientation map and circular variance.	90
6.11	Combined mask and reconstruction.	91
6.12	Examples of contrast image.	91
6.13	Candidate selection (hemorrhages).	92
6.14	Examples of final hemorrhage candidates.	93
6.15	Vessel likelihood feature.	94
6.16	Features importance (hemorrhages).	95
6.17	Classification result (hemorrhages).	96
6.18	FROC analysis (hemorrhages).	97
6.19	Examples of false negatives (hemorrhages).	98
6.20	Examples of false positives (hemorrhages).	99
6.21	ROC analysis and AUC (hemorrhages).	99
7.1	ROC curves on e-ophta database and recent database.	103
7.2	Incorrectly classified cases.	105
A.1	Examples of segmented FOV.	111
A.2	Example of bad quality image.	112
A.3	Background noise level differs between images.	112
A.4	Result of ASF.	113
A.5	Threshold defined by the largest flat zones.	113
A.6	Toggle-mapping is used to smooth the edge.	114
A.7	Result of FOV segmentation.	114
B.1	Examples of difficult cases for vessel segmentation.	118
B.2	Vessels segmentation based on alternating sequential filter.	119
B.3	Vessels segmentation results (ASF).	120
B.4	Gabor Kernel.	121
B.5	Results of Gabor filters.	121
B.6	Thresholding on the supremum of the results of Gabor filters.	122
B.7	Orientation information from Gabor filters.	123
B.8	Illustration of width features extraction.	124
B.9	Vessels segmentation results comparison.	125
B.10	Vessels segmentation results comparison.	126
B.11	Vessels segmentation results comparison.	127
C.1	Examples of difficult cases for OD detection.	130
C.2	Border reflection	131
C.3	Detection of border reflections.	131
C.4	Intensity analysis.	132
C.5	Skeleton of vessels mask.	133
C.6	Vessel projection to the horizontal axis.	133
C.7	Image without OD.	134
C.8	Estimation of OD position using vessel information.	135
C.9	Example of not detected OD.	135

LIST OF FIGURES

List of Tables

3.1	Comparison of scale coefficients.	33
4.1	Feature list	50
4.2	Details of pixel level validation.	54
4.3	Comparison of AUC of exudates detection on three public database.	57
5.1	Candidates selection.	68
5.2	Descriptor list	72
5.3	Sampling of FROC curve.	78
7.1	Confusion matrix of the results on e-optha database.	102
7.2	Confusion matrix of the results on the recent database.	104
C.1	Number of pixels in skeleton image in three directions.	134

LIST OF TABLES

Diabetic retinopathy is the main cause of blindness among the middle-aged population in developed countries. 347 million people (2008) worldwide suffer from diabetes (Danaei et al. [20]). A recent study shows that the general prevalence of diabetic retinopathy among diabetic patients is close to 28.7%, and yearly accumulative incidence of diabetic retinopathy is 2-6% (Delcourt et al. [23]). Moreover, after 15 years of diabetes history, 2% patients are blind and 10% patients suffer from poor vision (Massin and Erginay [65]). The first stage of the disease is silent, therefore a regular annual follow-up is recommended to diabetic patients. Moreover, given the increasing number of diabetic patients, and the limited, or even decreasing (in France), number of specialists, annual eye screening for patients is not performed sufficiently. Telemedicine networks largely improve patient follow-up, becoming a solution. The OPHDIAT telemedicine network for diabetic retinopathy screening (Erginay et al. [25], Massin et al. [67]) was established in 2004 by Assistance Publique - Hôpitaux de Paris (AP-HP). Statistics show that 75% of the exams in this network are considered as normal by the readers (Erginay et al. [25]). In this context, the TeleOphta¹ project aims at developing an automatic system for detecting normal exams (Decenci re et al. [22]).

1.1 TeleOphta Project

Annual diagnosis of diabetic retinopathy is insufficient in France due to various reasons. Lack of ophthalmologists and the increasing number of diabetic patients are two important factors. The OPHDIAT telemedicine network has been put into operation by Assistance Publique - Hôpitaux de Paris (AP-HP) since 2004, to improve diabetic retinopathy screening. The system structure is illustrated in Fig. 1.1. OPHDIAT consists mainly of three parts. Screening centers are responsible for taking digital fundus images, typically two per eye, and sending them to the servers, accompanied with patient information. They are also in charge of receiving the recommendations sent back by the ophthalmologist from the reading centers, and of communicating their advice to the patient. The servers transfer and store

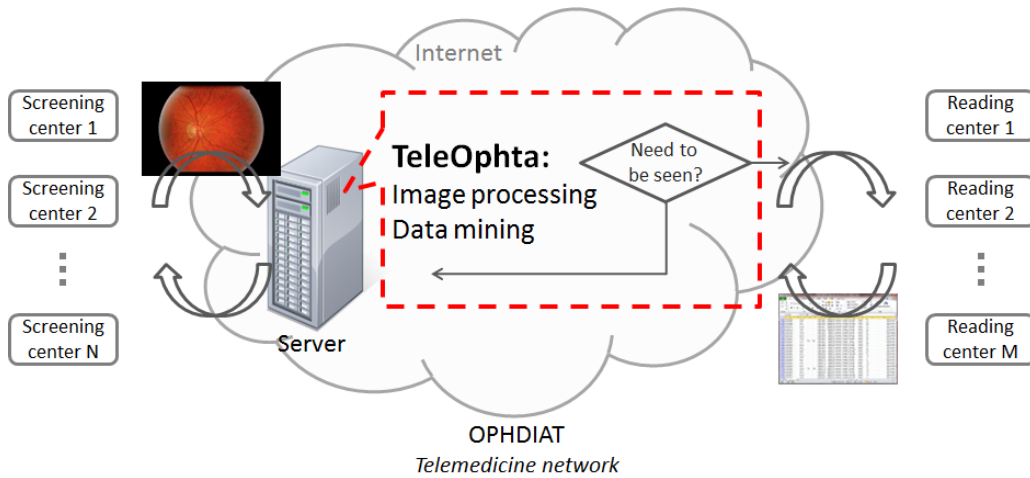


Figure 1.1: OPHDIAT telemedicine network structure and TeleOphta project.

the patients records and the reading results. Reading centers are where the images are reviewed by specialists.

In 2010, 34 screening centers were associated to OPHDIAT, and 13121 screening acts were performed. A telemedicine network like OPHDIAT largely optimizes the resource utilization rate. It simplifies the deployment of screening centers and diagnosis requirements (a computer with Internet connection is sufficient). However, it doesn't change the reality of lacking ophthalmologists. Moreover, their workload may become heavier, because of the simplified image capture conditions. On top of that, telemedicine acts are reimbursed by the French national health service since February 2014. New screening centers are expected to join the existing telemedicine networks, and new networks should be created.

Among the examinations performed within OPHDIAT every year, about 74% are considered healthy and 26% are referred to an ophthalmologist for either diabetic retinopathy or other pathologies (Massin et al. [67]). Computer aided diagnosis system can be an efficient way to solve this problem, automatically detecting healthy cases before transferring to the reading centers. Automatic methods to help reducing the burden on specialists are actively developed (Scotland et al. [90], Philip et al. [76], Abràmoff et al. [4], Dupas et al. [24], Agurto et al. [5]). In this context, the TeleOphta project was proposed aiming at developing a system to classify the examinations acquired in the network (Decencière et al. [22]).

TeleOphta is a collaborative research project, funded by the French National Agency of Research (ANR) through the TecSan program. The collaboration team includes the following partners:

- Assistance Publique - Hôpitaux de Paris (AP-HP), which provides medical expertise and anonymous clinical data.
- ADCIS, a software company specialized in image processing and analysis, responsible for the final integration of the developed methods.

¹Details of TeleOphta project refer to <http://teleophta.fr>

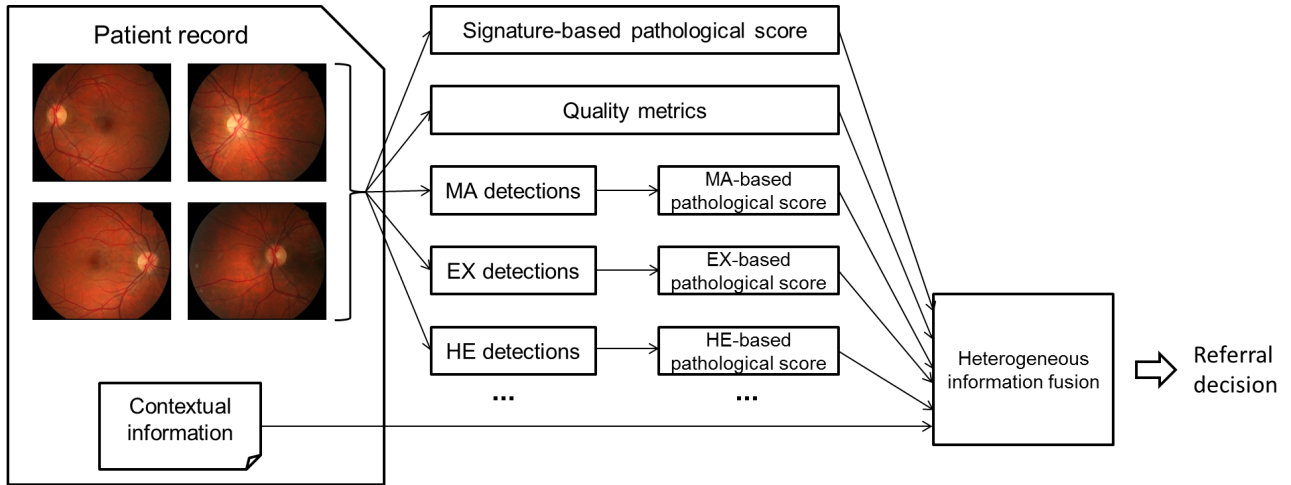


Figure 1.2: TeleOphta strategy. Workflow of Patient record classification.

- Laboratoire de Traitement de l'Information Médicale (LaTIM - UMR1101 INSERM, Telecom Bretagne), responsible for the development of statistical learning methods to classify examinations.
- ARMINES/MINES ParisTech (Centre de Morphologie Mathématique), project coordinator and in charge of the development of the image processing methods, such as retinal structures segmentation and lesions detection.

The main objective of TeleOphta is to develop a multidisciplinary research in automatic patient record classification by combining data mining and digital image processing concepts. Numerous studies on automatic fundus image processing have been done during the past 15 years. Meanwhile, with the development of new fundus image capture technologies, it still remains an active domain. Most studies concerned relatively small and homogeneous databases. However, TeleOphta is positioned in a real clinical environment and faces a large variety of images. To deal with this challenge, a large database with heterogeneous fundus images named e-ophta is dedicated to this research. The strategy of TeleOphta is illustrated in Fig. 1.2. For each patient record, image based information and contextual information are extracted. Machine learning methods are used to fuse heterogeneous information. Finally, each exam will be classified as “To be referred” or “Not to be referred” to a specialist.

1.2 Medical background

This thesis focuses on fundus image processing, so as to detect diabetic retinopathy lesions in the image. Minimum anatomical and medical context are presented in this section. Information in this section comes mainly from Massin et al. [66].

1.2.1 Retina and fundus images

The eye is located in the orbit, a cavity in the skull. It is connected to the brain via nerve fibers, which join in the optic nerve. The fundus of the eye is composed of three layers: sclera, choroid and retina

(see Fig. 1.3). The retina is located in the inner surface of the eye. It is a transparent and thin layer (less than 0.5mm of thickness) but is the most complex structure in the eye.




Figure 1.3: Anatomic image of the eye (from wikipedia) and fundus image.

An eye fundus image is taken by a “fundus camera”, which is a specialized microscope with an attached camera. A typical fundus camera views 30 to 50 degrees of retinal area, with a magnification of 2.5x. The observation light goes through a series of lenses, and enters the eye through the cornea onto the retina. The reflected light from the retina goes back to the microscope and the camera captures the image immediately. An example of fundus image is shown in Fig. 1.3.




Figure 1.4: Different capture fields. a) Macula centered image, b) Optic disc centered image, c) Temporal field of the macula.

Optic disc and macula are two important structures in the retina. Their anatomical positions and corresponding positions in fundus images are shown in Fig. 1.3. Blood vessels come into the retina through the optic disc. This point is also known as the blind spot, because it doesn't contain

photoreceptor cells. In eye fundus images, the optic disc appears as a white ellipse. In the direction from optic disc to the temple, we find the macula and its center, the fovea. This is the most light-sensitive area in the eye. The macula contains no vessels, and appears generally as a dark region. Optic disc and macula are two reference points in fundus images. Their position reveals basic information, for example, the left or right eye. Given the limited angle of fundus cameras, only a part of the retina is captured in each image. The diagnosis protocol of diabetic retinopathy states that the diagnosis should include at least two fundus images of 45° per eye, one macula centered, and another optic disc centered (Fig. 1.4). A third image covering the temporal field of the macula is optional (Massin and Erginay [65]).

The vascular network provides blood supplies to the retina. Blood vessels have lower reflectance compared to other retinal structures. Thus, they appear darker than the background in fundus images. Veins are darker than arteries. In fundus images, after entering the retina via the optic disc, the main vessels go vertically, and then bend towards the macula. Sometimes, vessels in the choroid layer are visible, probably due to a lack of pigmentation of the retinal pigment epithelium.

1.2.2 Diabetic retinopathy

Diabetic retinopathy is one of the top five causes leading to blindness, and is the first cause of blindness for people less than 50 years old (Massin et al. [66]). It is a complication of diabetes with a very complex pathogenesis. Most of the time, it has no early warning signs. It is a progressive microangiopathy characterized by small vessel damages and occlusions. In the early stage, also known as Non-Proliferative diabetic retinopathy (NPDR), the capillaries develop tiny dot-like outpouchings called micro-aneurysms, while the retinal veins become dilated, tortuous, and associated with multiple hemorrhages and intra-retinal exudates. In the advanced stage, also known as Proliferative diabetic retinopathy (PDR), retinal ischemia eventually stimulates the formation of fragile new vessels. New vessels can cause the vitreous to separate from retina. If the vessels bleed, a massive hemorrhage may cause sudden visual loss.

The detection and diagnosis of diabetic retinopathy in the early stage can help to tremendously slow the degeneration. At this stage of the disease, lesions such as microaneurysms and hemorrhages are likely to be present. Thus an automatic detection of these lesions can help to make the diagnosis of diabetic retinopathy easier, better and more reliable. Three main lesions are concerned in this work, and presented in the following:

- Microaneurysms: dilations of the venous end of retinal capillaries, of variable size, mostly between 10 and $100\mu m$, but not above $125\mu m$. In color retinal images, they appear like little dark red dots (or dark dots in the green channel) detached from blood vessels. They are the first sign of diabetic retinopathy (see Figure 1.5(a)).
- Hemorrhages: blood leaks within the retina. They can appear anywhere in the retina, with any size and shape. There are many kinds of hemorrhages such as dot hemorrhages, blot hemorrhages, flame hemorrhages. In color retinal images, hemorrhages appear like dark red regions (see Figure 1.5(b)). The smallest hemorrhages are very similar to microaneurysms.
- Exudates: accumulations of lipidic deposits within the retina. They appear yellow in color retinal images (or as bright regions in the green channel) (see Figure 1.5(c)).

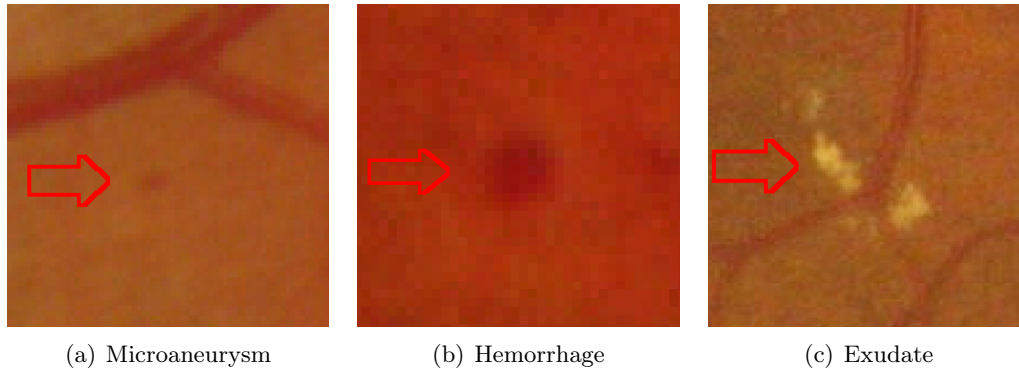


Figure 1.5: Clinical signs of diabetic retinopathy

1.3 e-ophtha database

The e-ophtha database was constructed by extracting and anonymizing the examinations from the OPHDIAT network during two years starting from 2008. It contains 25702 examinations. Each examination contains at least four images (two for each eye), and contextual information such as age, weight, diabetes history, etc. Annotations as diabetic retinopathy level and text comments are given by a specialist for each examination. All images are in compressed JPEG format, so as to speed up Internet transfer. The entire database occupies 50 GB of storage.

e-ophtha contains clinical data, acquired in clinical conditions, rather than a well-defined digital image processing dedicated database. The images are taken in different screening centers with different types of cameras. The images show important size, color, illumination, contrast and quality diversities (Fig. 1.6). Currently, 8 different image sizes are present in the database, ranging from 1440×960 to 3504×2336 pixels. This variation makes the image processing algorithms development more challenging. In addition, the block effect of JPEG compression cannot be ignored while dealing with small objects, for example, microaneurysms. On the other hand, a large and diverse database like e-ophtha with annotations is a precious resource for the study of fundus image processing and diabetic retinopathy auto-grading.

Three subsets, e-ophtha MA, e-ophtha EX¹ and e-ophtha HE have been created in the interest of lesion detection methods development. They respectively deal with microaneurysms, exudates and hemorrhages. Thanks to the annotation software developed by ADCIS (Fig. 1.7), lesions are meticulously and manually annotated by an ophthalmologist and confirmed by another.

The e-ophtha MA database contains, on one hand, 148 pathological images, on which 1306 microaneurysms and small hemorrhages have been annotated and, on the other hand, 233 “healthy” images (i.e. not containing microaneurysms).

The e-ophtha EX database contains 47 “healthy” and 35 pathological images. On the pathological images, 2278 exudate components have been precisely annotated. This number might seem very high. This is due to the fact that exudates typically appear as a constellation of small components, which have been individually annotated (see Fig. 1.7(b)). Note that many healthy images in the e-ophtha EX

¹Details and download of e-ophtha MA and e-ophtha EX database refer to <http://www.adcis.net/en/Download-Third-Party/E-Ophtha.html>




Figure 1.6: Fundus images examples from the e-ophtha database. These images show different sizes and colors, bright border reflections, non-uniform illumination, blur effects, etc.

database have been chosen because they contain bright normal structures, such as reflections, which might mislead automatic exudate detection methods.

The e-ophtha HE database contains 90 pathological images, on which 299 hemorrhages have been annotated. It shares the same 233 “healthy” images with the e-ophtha MA database, because these images are all dark lesions (microaneurysms, hemorrhages, etc.) free.

As a complement of the original annotation of the entire e-ophtha database, a subset of 500 images have been randomly extracted for a double reading. Thus, three expert opinions are available for this database, which is named e-ophtha DR (double-reading).

The e-ophtha database and its subsets: e-ophtha DR, e-ophtha MA, e-ophtha EX and e-ophtha HE, together with their annotations, and its diversity as used in real clinical condition, constitute a valuable scientific research basis. They are available on the website: <http://www.adcis.net/en/Download-Third-Party/E-Ophtha.html>

1.4 Research proposal

The main objective of this work is to develop image processing methods to detect diabetic retinopathy lesions in fundus images in a clinical context. Three important lesions are considered: microaneurysms, exudates and hemorrhages. Segmentations of main retinal structures are also necessary. The document is organized as follows:

Chapter 2 discusses morphological residues analysis. A new residual operator is proposed for the pre-processing of fundus images. It can separate bright and dark parts of the image with great details

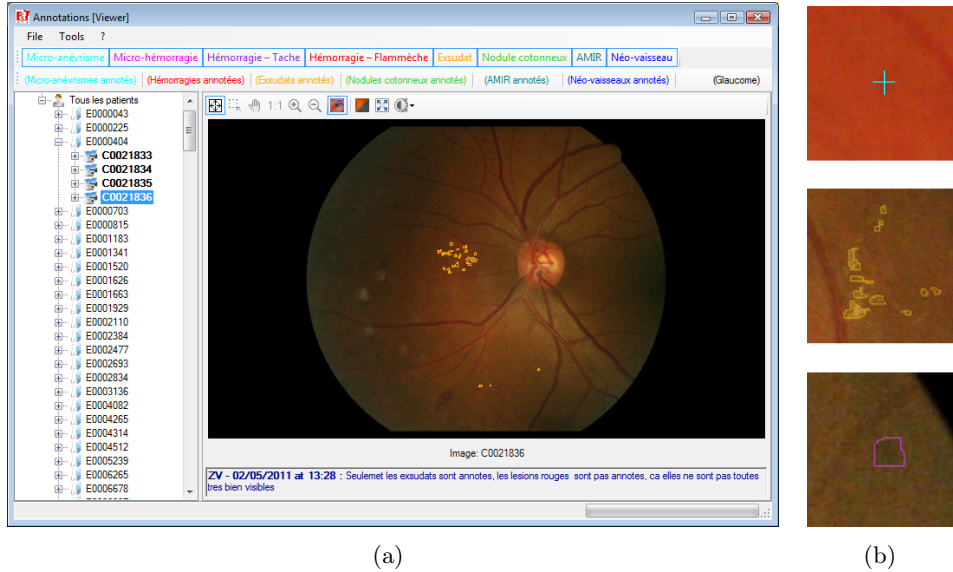


Figure 1.7: Screenshot of TeleOphta annotation software. a) Interface in viewer mode. Annotations are overlaid on the original image, b) Examples of expert annotations for microaneurysm, exudate and hemorrhage.

and little noise. Objects in fundus images generally have a clear relative distinction in intensity, for example, vessels are darker than its background, while exudates are brighter. Thus, the proposed operator is very useful for preliminarily extraction of objects of interest.

Chapter 3 introduces a new spatial calibration method for fundus images. Due to the diversity of the e-ophta database, and also under real clinical conditions, we need to deal with images with different resolutions. A general and efficient spatial calibration method is proposed.

Chapters 4 to 6 concern the detection of exudates, microaneurysms and hemorrhages. New methods are proposed and evaluated on e-ophta, as well as other public databases. Thanks to the ground-truth provided by the e-ophta databases, machine learning methods are also used in these parts to give an estimation of the probability of the detected lesions.

Chapter 7 presents the patient classification strategy, which gives the final probability as “To be referred” or “Not to be referred”. The result of TeleOphta evaluation are then presented and analyzed, and compared to human specialist.

Chapter 8 concludes the work and gives perspectives.

Annexes contain the methods for detecting and segmenting the main retinal structures. Because they are not the core interest of this work, but are indispensable, the methods will be briefly introduced in the Annexes.

Gray-level normalization is a basic problem to be solved in the first stage of image processing. Moreover, features extraction is an important step for object recognition or classification in computer vision, which is often in the last stage of image processing. However, the quality of extracted features is somehow affected by the gray-level normalization result. In this chapter, we will solve two questions: How to correct the non-uniform illumination; how to find a good representation for features extraction. A residue operator by alternating sequential filter and multi-threshold decomposition feature analysis method are proposed as solution. We will also discuss the application of multi-threshold decomposition to the extracted features in machine learning.

2.1 Pre-processing and residue operator

2.1.1 State of the art

Pre-processing is necessary for fundus images, because the contrast tends to diminish towards the edge of the images and the intensity varies between images. The color of fundus images is largely variable from person to person and between different ethnic origins. Moreover, the illumination within an image and between images can be significantly different. Many algorithms can largely improve the visual perception of the images. Histogram equalization (Youssif et al. [113]) is the most widely used. An improved version, adaptive histogram equalization (Wu et al. [111], Zuiderveld [122]) provides a better result. Adaptive local contrast enhancement (Sinthanayothin et al. [94], Walter et al. [109]) is another similar method which is based on a non-linear histogram transformation. Hoover and Goldbaum [48] and Foracchia et al. [33] suppose that the image consists of the original image with multiplicative and additive components. By means of foreground and background analysis, the average intensity and contrast are normalized to assigned values. The same model is used in Joshi and Sivaswamy [51], but a non-uniform sampling is used to estimate the components. Other methods such as contourlet transform are used in Feng et al. [26]. Complex methods involving retinal structures such as vessels removal (Ravishankar et al. [83]) are not considered in this stage, though they could give more promising results.

In the work of Antal and Hajdu [8], a number of pre-processing methods are compared in order to give a good microaneurysm detection. The comparison is based on counting the number of detected microaneurysms by implementing a general detection algorithm after the pre-processing. Their study shows that the method using contrast limited adaptive histogram equalization (Zuiderveld [122]), intensity adjustment (Gonzalez et al. [42]) and vessel removal (Ravishankar et al. [83]) give the best results. But the number of false positives is not considered in the evaluation. The vessel removal method, as explained, is not considered in this stage. The intensity adjustment method used by the author lacks a clear reference, and is in fact no more than a global histogram stretching method. Thus, we implemented the method of Zuiderveld [122] and another adaptive local contrast enhancement method from Walter et al. [109], which has a relatively mediocre performance according to the article, for comparison. A difficult image is chosen for this test (Fig. 2.1(a)). A bright region in the left border side and a dark region in the right side appears in this image. In fact, a linear histogram equalization is already applied on Fig. 2.1(a) for a better visualization.

The contrast limited adaptive histogram equalization (CLAHE) is often used in medical image processing. Compared to the classical histogram equalization methods, it is applied locally and limits the amplification by clipping the histogram at a predefined value. For more details, refer to the original article (Zuiderveld [122]). The result is shown in Fig. 2.1(b). The contrast is enhanced, but the illumination is poorly balanced. This method visually improves the image quality.

Walter et al. [109] defined a polynomial transformation to stretch the histogram within a moving window.

$$u = \begin{cases} \frac{\frac{1}{2}(u_{max}-u_{min})}{(\mu_f-t_{min})^r} \cdot (t-t_{min})^r + u_{min}, & \text{if } t \leq \mu_f, \\ -\frac{\frac{1}{2}(u_{max}-u_{min})}{(\mu_f-t_{max})^r} \cdot (t-t_{max})^r + u_{max}, & \text{if } t > \mu_f, \end{cases} \quad (2.1)$$

where t_{min} and t_{max} are the minimum and maximum values in the moving window, u_{min} and u_{max} are the expected output minimum and maximum values, μ_f if the mean value within the window and r is the polynomial parameter. For the result in Fig. 2.1(c), the window size is 50, $r = 3$, $u_{min} = 0$ and $u_{max} = 255$. The black background is gray because there are light fluctuations, which can be easily cleaned up. The contrast is enhanced. The illumination is well balanced between the bright zone in the left and dark zone in the right. However, noise in the dark zone is amplified, too.

For certain methods, the pre-processing does not really improve automatic segmentation or detection. Sometimes, they can correct the non-uniform illumination in the image, but with the side effect of amplifying noise. Another kind of pre-processing techniques could remove the noise, for example, a median filter. But, it is risky, as it may compromise microaneurysm detection. The idea of background and foreground analysis is interesting. After obtaining the partition, the histogram can be redistributed according to the partition, thus to correct the intensity distribution. Based on this concept, we will propose below a new method based on residue operators. It will not only be used for microaneurysms detection, but also be used for retinal structure segmentation. A comparison between the proposed method and the implemented methods from the literature will prove the advantages of our method.

2.1.2 Residue operator by alternating sequential filter

The foreground usually refers to the regions of the image containing retinal structures and lesions (bright and dark). The background corresponds to the rest, which contains no useful information. In most background and foreground analysis methods, for example Walter et al. [109], the background is set to 128 and the maximum and the minimum value of foreground to 255 and 0. The foreground




Figure 2.1: Pre-processing algorithms. a) Original green channel, b) Contrast limited adaptive histogram equalization (Zuiderveld [122]), c) Contrast enhancement by polynomial transformation (Walter et al. [109]), d) Negative residues of alternating sequential filter.

is extracted by a local histogram or contrast analysis. We refine the definition of the foreground separately for dark structures and bright structures. Morphological operators will be used to extract the foreground. As a result, the background becomes 0 and the foreground always has a positive value greater than the background. We will focus on the dark structures extraction, while the bright structures could be obtained by the dual operation.

Sometimes, it is easier to remove interesting objects than to remove all irrelevant objects. Residue operators such as the Top-Hat transforms are morphological operators used to extract interesting objects. By considering the size or shape of the target object, a simple morphological operator such as an opening can be used to remove it. Then, they are recovered by an arithmetic difference. The Top-Hat was already used in this domain for vessel segmentation in the work of Aquino Martín [9]. Walter et al. [109] used the residue of closings by diameter to get microaneurysm candidates. Other advanced residue operator, such as ultimate openings, which is a multi-scale morphological operator, introduced by Beucher [11], may be used to extract the most contrasted structures and the corresponding size

information from an image. It is used in this work too, for the first time in this domain. Details will be introduced in the following chapter.

In our case, in order to get dark structures which normally have a limited size (width of the vessels), a Top-Hat by closing (hexagonal structuring element with size 6 in the following example) gives a good result. An example of an image containing vessels and microaneurysms is shown in Fig. 2.2(a), and the result of the Top-Hat in Fig. 2.2(b). The remaining non-zero pixels belong to the foreground. It can be seen in this figure that a single morphological closing with a fixed size keeps too much background variations.




Figure 2.2: Residues operators. a) original image, b) Top-Hat by closing, c) Negative residues of opening and closing (Top-Hat extension), d) Negative residues of ASF, e) Negative residues of ASF by reconstruction.

In order to improve the classical Top-Hat, Salembier [84] proposed an extension to the Top-Hat. Let γ be an opening, φ the dual closing and I an image. The classical Top-Hat τ is:

$$\tau(I) = I - \gamma(I) \quad (2.2)$$

The Top-Hat extension ρ is:

$$\rho(I) = I - \{\gamma\varphi(I) \wedge I\} \quad (2.3)$$

The result is shown in Fig. 2.2(c). The background is very clean, and the two microaneurysms are kept. However, small vessels are erased.

We propose to use a multiscale approach. The alternating sequential filter (ASF) is a morphological operator combining openings and closings with increasing sizes of structuring elements (Sternberg [100], Serra [92] and Soille [95]). Denote γ_i an opening of size i and φ_i a closing of size i . $\{\gamma_i\}$ and $\{\varphi_i\}$ have to be respectively a granulometry and an anti-granulometry. Two families of alternate filters can be defined in the following way:

$$\begin{aligned} m_i &= \gamma_i \varphi_i, \\ n_i &= \varphi_i \gamma_i. \end{aligned} \tag{2.4}$$

Finally, two ASFs of size i are defined by:

$$\begin{aligned} M_i &= m_i \cdots m_2 m_1, \\ N_i &= n_i \cdots n_2 n_1. \end{aligned} \tag{2.5}$$

ASFs are usually used to reduce noise in images and are especially interesting for removing multi-scale noise while preserving large objects. In our case, we use it in a reverse way. The negative residue of the ASF is what we are interested in. A new Top-hat extension ρ_{asf} is defined by:

$$\rho_{asf}(I) = \{M_i(I) \vee I\} - I \tag{2.6}$$

The result on our test image is shown in Fig. 2.2(d) (the size of the filter is set to the estimated vessel width). It is clear that the background noise has been reduced. The result on the entire image for the previous example is shown in Fig. 2.1(d).

If γ_i and φ_i are replaced by the reconstruction opening γ_i^{rec} and reconstruction closing φ_i^{rec} , the corresponding ASFs become M_i^{rec} and N_i^{rec} . Thus, replacing $\gamma\varphi$ in Eq. 2.3 by M_i^{rec} gives a new Top-Hat extension. The result of the application on the previous image is shown in Fig. 2.2(e). We can see that small vessels are not kept in the residue, compared to the negative residue of ASF. This is because the reconstruction closing preserves the small vessels connected to larger vessels, which are removed by the morphological closing. Thus they don't appear in the residues. As a residue operator, the classical ASF is better than the reconstruction ASF for our application. The classical Top-Hat and the proposed method are applied on other images with the same structuring element size (Fig. 2.3). The residues by ASF generally appear cleaner than the classical Top-Hat.

Back to the visual comparison between pre-processing methods (section 2.1.1), the results of the state-of-the-art methods (Fig. 2.1) are manually thresholded. The thresholding ranges are quite different for each image, while the objective is to segment dark structures, especially blood vessels, and keeping as less noise as possible. The thresholding result of CLAHE (Fig. 2.4(a)) keeps most large vessels, but also a large zone on the right side. This is the major drawback of contrast enhancement algorithms: they tend to amplify noise. The thresholding result of Walter's method (Fig. 2.4(b)) shows less noise on the right side and more vessel details. However, in the red rectangular region, where a vessel with bifurcation should appear, the threshold is not enough to segment this vessel (Fig. 2.4(d)). Fig. 2.4(c) is the result of thresholding the result of ASF. Most vessels are segmented with reduced noise on the right side. The same bifurcation vessel is more complete (Fig. 2.4(f)). But in the green window, the large vessels in the dark zone have disappeared with the noise. This may be a drawback for vessel segmentation. As for lesion detection, it is always more favorable with less noise. However, although the vessels appear in the result of Walter's method (Fig. 2.4(e)), there is much noise, which makes them indistinguishable by a simple threshold. For the proposed method, the original contrast is kept. Reducing the threshold will get more vessels and noise. Since the starting level is set to zero, it is easy to guarantee the segmentation quality of the major structures, while treating the noisy regions




Figure 2.3: Residue operators applied on other images. a) and d) original images, b) and e) Top-Hat by closing, c) and f) Negative residues of ASF.

differently. Generally, the image is still noisy but it is almost uniformly distributed, which means that illumination is corrected. Compared to the other algorithms, the proposed one preserves more details. The reasons why CLAHE is ranked first in the cited work could be various: their evaluation strategy is not totally satisfactory; their candidate extractor is specific to their images; their database does not contain complex images.

The advantages of the proposed method are manifold:

- It directly works on the original image and takes only one parameter - the maximum size of the structuring element, which is easy to tune.
- ASF is a multi-scale filter. It preserves the details to a great extent while removing background variations.
- Dark and bright structures can be obtained separately by dual operators. The pre-processed image contains only the desired parts (bright or dark), which reduces the interferences.
- The contrast is normalized by taking residues. The starting level of all dark structures is set to zero, which facilitates the following steps, for example, thresholding.




Figure 2.4: Thresholding of a pre-processed image: a) Thresholding from 23 to 49 on the result of CLAHE (Zuiderveld [122]); b) Thresholding from 10 to 88 on the result of the contrast enhancement (Walter et al. [109]); c) Thresholding from 5 to maximum value on the result of the ASF; d), e), f) and g) Details.

ASF are computationally intensive when compared to a single closing or a histogram normalization, which is a drawback. However, efficient algorithms and implementation of morphological operations make the computing time compatible with our application.

2.2 Multi-threshold decomposition

The resulting pre-processed image not only facilitates thresholding, it also provides more useful and reliable features. It normalizes intensity related features inter-images and intra-images. It also largely facilitates the interest of threshold based decomposition. In this section, we will discuss the question about “How to find a good representation for the features”, based on the proposed pre-processing method.

2.2.1 Different threshold choice strategies

The selection of the features and the computation support vary between different contexts. In our case, the objects need clear borders for features analysis, because some objects are very small and sensitive to size and geometry variations. For example, Fig. 2.5(a) shows a microaneurysm in the middle of the image. If we try to get its binary mask by thresholding, we will obtain several possible results shown in

Fig. 2.5(b) - Fig. 2.5(d). It is hard to tell which one is the best to represent this microaneurysm, because it depends on the features and classifiers. But the last mask is two times larger than the first one. Moreover, their length, circularity, average contrast, etc. change with different threshold values. Thus, for the same object, different borders could introduce different feature values and different classification results.




Figure 2.5: Thresholding of a microaneurysm. a) Original image (contrast enhanced), b) Thresholding under 72, c) Thresholding under 77, d) Thresholding under 82.

In the literature, we can find similar discussions on connected filtering and segmentation methods. In the work of Jones [50], the author classified the connected filters applied on gray-level images into two classes: flat component filters (Heijmans [46]) and nonflat ones. A component tree image representation and an attribute signature were proposed. Nonflat filters can link threshold components at sequential gray-levels in images. The sequential analysis discloses how does the component change with the thresholds. Back to the problem raised in the beginning of this section (Fig. 2.5), this analysis is no longer performed on one threshold, but between a range of gray-levels. A similar approach can be found in the work of Pesaresi and Benediktsson [75], where an analysis based on the derivative of the morphological profile was proposed. The residue analysis ultimate opening proposed by Beucher [11] uses also the “vertical” information between different thresholds. The recently proposed methods based on differential area profiles by Ouzounis et al. [74] and the segmentation method based on attribute profiles and area stability by Serna et al. [91] are also very interesting.

The state of the art shows that useful information can be obtained by a sequential analysis between thresholds. They can be applied to gray-level images to extract synthetical features. For example, the ultimate opening is used in the following as a candidate feature. However, with the help of the proposed pre-processing method, non-uniform illumination is corrected. Features are computed on the residue image. Therefore, another approach consists in directly using the primary features computed on each threshold. The threshold choice strategy depends on different objects, as we will see below.

For objects which have no specific size, shape or contrast, such as exudates (Fig. 2.6(a)), we propose to start from the bottom layer (white region in Fig. 2.6(b)). The bottom layer means the minimum threshold for obtaining the binary mask. Then, several layers will be sampled uniformly between the maximum and minimum intensity values by thresholding. The number of layers depends on the requirements and the computation load. For example, a two-layer strategy is adopted for certain exudate features extraction (Fig. 2.6(b)).




Figure 2.6: Multi-threshold on an exudate. a) Image after pre-processing, b) Thresholding at bottom (white) and middle (red) layers.

For objects like microaneurysms, which have specified shape and limited size and contrast, we propose to start from the top. The maxima of each connected component is the starting point for feature calculation. Then the same feature will be computed on the following layers by decreasing the threshold.




Figure 2.7: Max-tree decomposition and node selection (Original profile shown in Fig. 2.8).

In practice, max-tree image decomposition (Salembier et al. [85]) is implemented to speed up this calculation and select the components. Each node C_h^k in the max-tree contains pixels in the k th connected component at intensity level h . An illustration is given in Fig. 2.7 using the profile structure in Fig. 2.8(a). Two similar profiles are shown in Fig. 2.8(a): A microaneurysm (left) and another structure (right). Let us use the second structure to illustrate the max-tree decomposition and node selection. Each node points to its father node. The feature calculation starts from the maximum, which is node C_8^0 in this example, then passes down to its father node (C_4^0). This process is repeated until the root node is reached. Note that nodes C_5^0 and C_4^1 are dropped (red crosses in Fig. 2.8(a) and Fig. 2.7) during the computation. Since microaneurysms can be roughly approximated by a Gaussian

shape (5), the selection of the nodes can remove the branches and compare only between the main parts of the structures. It also simplifies the computation.

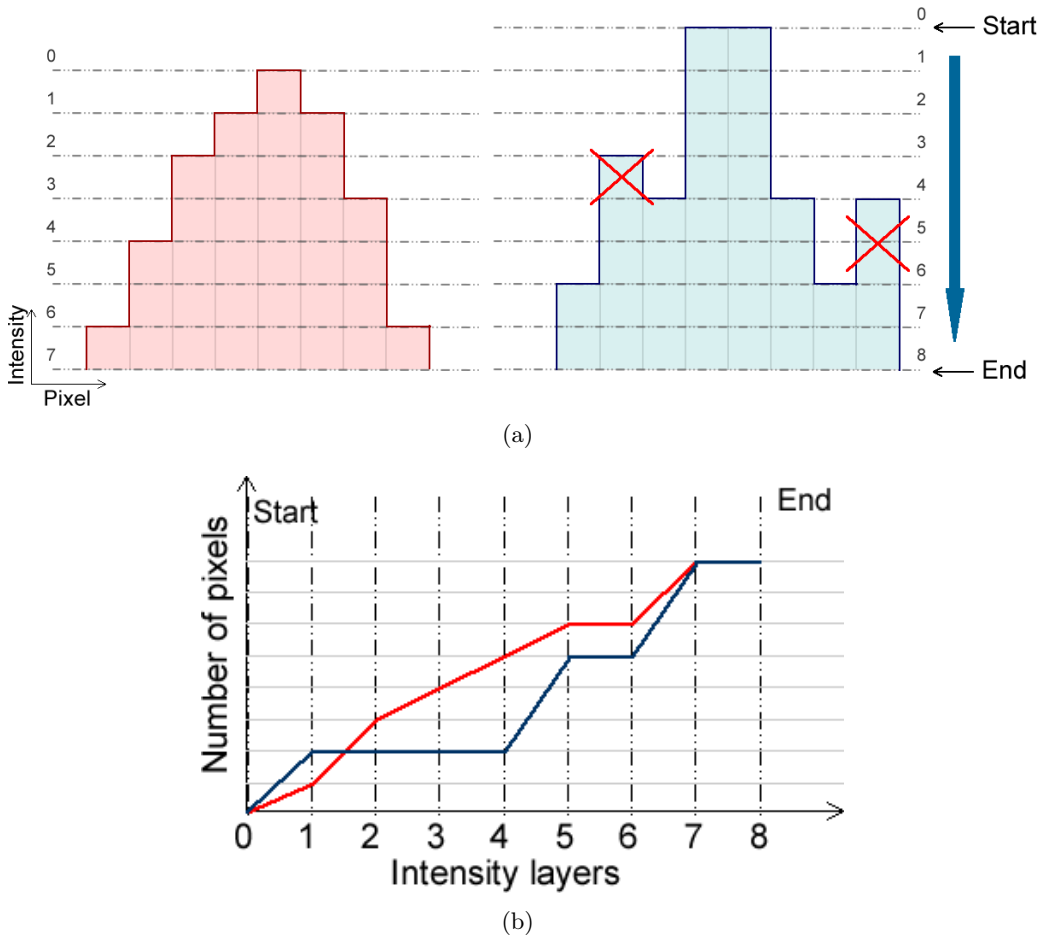


Figure 2.8: Layer width as a function of the Max-tree layer. a) Two similar profiles have the same intensity level but different Max-tree layers. b) Width evolution with respect to intensity level.

In order to distinguish the two structure profiles in Fig. 2.8(a), the width feature is computed. However, a general width or contrast analysis based on the bottom layer or the middle layer will not differentiate them. But if we represent their evolution along with the intensity layers, different curves are obtained. Fig. 2.8(b) shows the evolution of the width with respect to the intensity. The blue curve rises slower than the red one and has a jump at the middle, because of its abrupt width increasing in the middle. Note that the alignment is from the top level. No matter how high the contrast is, the features are always noted from the maxima. Considering the individual contrast and size variation, the alignment from the top has smaller variation than the alignment from the bottom layer.

In a practical way, we limit the number of layers during the computation. A statistical analysis shows that the maximum intensity level of microaneurysms in the residue image are mostly below 20 (see Chapter 5). Under this condition, the number of layers is set to 20. For the candidates whose contrast is less than 20, the rest of the features are automatically filled with an infinite value.

2.2.2 Machine learning application

Machine learning is widely used in image processing to identify objects. Two important parts should be considered while designing a machine learning process. First, the features extraction. Features should be representative and distinctive. Second, the choice of the machine learning method. We will concentrate on the first question in this part.

Features can be computed on the whole image, for example, the global mean and variation, or the coefficients of pyramid or wavelet analysis. Features can also be computed on connected components. This is the most common way for object segmentation, which is also widely used in our case. After obtaining candidates (connected components) or regions of interest, features like area, length and mean intensity etc. can be computed for each of them. Another common way is to extract pixel-wise features such as different color components. When introducing multi-threshold decomposition, the discussion is based on the features extracted from the connected components. However, by multi-threshold decomposition, some features are not only computed on the whole candidate, but also on a smaller support, which corresponds to a different level of the candidates contrast image. We have illustrated that in some cases considering different levels brings useful information into the classification. The normalization of the features is important, especially considering a heterogeneous database like e-optha. Thanks to the residue analysis, the contrast features computed on the residue image has already been normalized. The size will be normalized by spatial calibration (Chapter 3).

In our case, a Random Forest (Breiman [13]) is used to perform the classification of lesion candidates, such as microaneurysms and exudates. It has two important parameters to be defined: the number of trees and the number of features to consider for a best split. According to the author, random forests do not overfit. After a certain threshold, increasing the number of trees will not bring any improvement, but can reduce the variance of the result. In exchange, the computing time increases. We find that 500 trees provide a good and stable classification in our case. The number of features to consider for a best split is set to the square root of the number of total features, as recommended by Breiman. Moreover, Random Forests can evaluate the importance of features during the generation of the forest. Thus, a selection of the features is possible and easy to achieve. Details of the configuration and the results of Random forest will be discussed in the following chapters for each special purpose.

2.3 Conclusions and perspectives

In this chapter, we have first proposed using alternating sequential filters for residue analysis. The residues keep details and correct non-uniform illumination while removing background variations. The separation of dark and bright parts can reduce the interferences. This largely facilitates thresholding and contrast analysis. However, in extremely low contrast regions, the structures are lost together with the noise. Based on the above advantages, we have proposed a multi-threshold decomposition analysis method. It allows to analyze features based on connected components on different levels, which provides richer information than the analysis on the whole object. Different alignments of the features are introduced. A bottom-top alignment is appropriate for large structures and a top-bottom alignment is suitable for small objects. The applications in the following chapters for lesions detection will prove its efficiency.

Several points remain to be discussed. While extracting features on a Max-tree decomposition, branches which are not inherited from the highest maxima are pruned. For microaneurysms, which can be modeled by a Gaussian shape and have small size, pruning isolated nodes could simplify the model and accelerate the computation. However, for large structures as exudates, these pruned nodes may

CHAPTER 2. RESIDUE ANALYSIS METHODS

contain texture information. A difference between the computation with and without node pruning may reveal such information. The number or the volume of the pruned nodes may also provide interesting information.

Spatial calibration of fundus image

The heterogeneity of clinical retinal image databases, regarding image scale, contrast and quality, makes the design of generic image processing algorithms difficult. In this chapter, we focus on the spatial calibration of images. Two images from the e-ophtha database are shown in Fig. 3.1. They are respectively among the smallest and the largest images of the database. The width of the largest vessels measured in pixels, shows a large difference. How to solve this diversity and to make the following algorithms easily parameterized, is the main objective of this chapter. The spatial calibration is the base for other operations. At first, different possibilities of spatial calibration of fundus images are analyzed. Then a simple and efficient method which uses the width of the field-of-view (FOV, the visible round region in fundus images) as invariant is proposed (Zhang et al. [119]). It will be validated on all image resolutions from the e-ophtha database.




Figure 3.1: Images of different resolutions. a) Image size 1440×960 pixels; main vessel width is equal to 9 pixels, b) Image size 3504×2336 pixels; main vessel width is equal to 28 pixels.

3.1 Possible approaches

In the literature, most studies use a small (usually less than 100 images) and homogeneous database to evaluate their algorithms. They are not concerned by image calibration. Some works deal with a large database, such as Fleming et al. [30] and Hipwell et al. [47]. Thousands images are used to evaluate their methods, but the images were captured by the same camera or have the same size.

Tang et al. [102] used Messidor database (Dupas et al. [24]) which contains 1200 images with 3 types of resolution, from 1440×960 to 2304×1536 . The FOV was detected automatically and the images were rescaled to 1026×681 pixels to have a FOV width approximately equal to 630 pixels in diameter. The authors didn't provide details. The limitation of this method is that it rescales the images by the image size, though the FOV width is mentioned. Another similar method was used in the work of Fleming et al. [29]. A large database containing images of 45° with resolutions from 1024×1536 to 1440×2160 was used in their work. According to the author, images were scaled so that the vertical dimension of the visible fundus was 1400 pixels. Since details were not provided by the author, we interpret "the visible fundus" as the FOV, while the "vertical dimension" should be the horizontal dimension in normal cases (the images' height is larger than the width according to the author, which is normally the contrary). Using the width of the FOV to rescale images is more appropriate. The same method can be found in the work of Sánchez et al. [88]. Their database contains 144 images captured by three cameras with two capture angle 35° and 45° . The resolutions vary from 768×576 to 2048×1536 . The images were automatically resized to have a FOV with a standardized diameter of 650 pixels. The problem is that their images have two different capture angles. Using the width of the FOV under this condition, will introduce a bias after spatial calibration. The reason will be discussed in the following sections.

The most common method consists in using the image size directly to perform the normalization. Before any processing, all images are rescaled to the same size (usually to the same width). However the image size cannot precisely represent the pixel size. For the same example in Fig. 3.1, if we calculate the ratio between the vessel width, we get 3.1 (28/9). However, the ratio between the image widths is 2.4. The bias will be further increased for surface measures. The size of the FOV, is a better reference for spatial calibration, especially the width of the FOV, because they have different shapes: round disc (Fig. 3.1(a)) or truncated disc (Fig. 3.1(b)). The ratio between the FOV widths is 3 (2656/884).

Spatial calibration based on anatomic structures and lesions size, such as the diameter of the optic disc (OD) and the width of blood vessels, is another possible approach. However several difficulties make it hard to realize. Firstly, the anatomical size of structures varies from one person to the other (for example, the theoretical size of the OD is between 1 and 2 mm). Secondly, obtaining a precise segmentation and measurement of retinal structures, such as blood vessels, remains a problem.

The distance between OD and Fovea could be used to normalize images (Fig. 3.2). Indeed, this distance is relatively constant in human eyes. In fact, the diameter of the human eye shows small variations between individuals. But the problem is that the OD and Fovea need to be detected, which is not easy, and moreover, they may be absent from some images.

3.2 Spatial calibration using the width of the field-of-view

The FOV width can be effectively used to perform the spatial calibration. But in the literature, it has never been systematically discussed (to the limit of our knowledge). In this section, we will explain how it works and therefore propose a new spatial calibration method without re-sampling the image.




Figure 3.2: Normalization based on the distance between the OD and the Fovea.

The human eye reaches its adult diameter of $24mm$ at age 13. Size variations between individuals are typically smaller than $1mm$. Given that in most diabetic screening program, the capture angle is constant, the diameter of the FOV in retinal images can be considered as an invariant (illustrated in Fig. 3.3). For example in the case of OPHDIAT, the capture angle is always equal to 45° . However, between different capture angles, the transition needs to be studied. This is the reason why the method used in Sánchez et al. [88], which also uses the width of the FOV but ignores the capture angle, is not appropriate.




Figure 3.3: Capture angle of fundus image.

The FOV is generally easy to segment, though there exist several extreme cases in the e-optha database, which contain a very dark zone (Fig. 3.4). A method is proposed to segment the precise




Figure 3.4: Image contains a dark zone and the segmented FOV.

contour of FOV and is detailed in Appendix A. This method is efficient. It can however be replaced by any other FOV segmentation method.

The spatial normalization is accomplished by giving estimations of the sizes for main retinal structures and lesions according to the width of the FOV. Thus, there is no need to resize the images and losing original information. Four useful sizes will be used in the following chapters: the width of the FOV (d_{FOV}), the height of the Optic Disc (d_{OD}), the maximum vessels width (d_{vessel}) and the maximum width of microaneurysms (d_{MA}). They are illustrated in Fig. 3.5. Other lesions such as exudates and hemorrhages are not included because they could have any size. The four presented structures and lesion sizes are generally definitive.




Figure 3.5: Spatial normalization by estimations of the sizes of retinal structures and lesions.

The diameter of the FOV, d_{FOV} , can be easily obtained by measuring the width of the segmented FOV. Then the other sizes are given by:

$$d_{OD} = \frac{1}{5}d_{FOV} \quad (3.1)$$

Image size	Distance between OD and Fovea	SC_{Manu}	Width of FOV	SC_{Auto}	$SC_{Classic}$
1440 × 960	349	1	910	1	1
1504 × 1000	352	0.99	923	0.99	0.96
1728 × 1152	512	0.68	1305	0.70	0.83
1936 × 1288	555	0.63	1484	0.61	0.74
2048 × 1360	595	0.59	1533	0.59	0.70
2544 × 1696	778	0.45	1924	0.47	0.57
2816 × 1880	817	0.43	2159	0.42	0.51
3504 × 2336	1044	0.33	2650	0.34	0.41

Table 3.1: Comparison of scale coefficients between the manual measurement (SC_{Manu}), the proposed method (SC_{Auto}) and the classical method (SC_{Class} , using image width).

$$d_{vessel} = \frac{1}{74}d_{FOV} \quad (3.2)$$

$$d_{MA} = \frac{1}{99}d_{FOV} \quad (3.3)$$

These values are empirically chosen and validated in the applications. These estimations will be used to parameterize other algorithms. In certain particular circumstances, images are resized to facilitate and accelerate the calculation. For example, while localizing the OD, image details are not needed. Then d_{FOV} will be used as the invariant. Images are resized according to the d_{FOV} rather than the image size.

3.3 Validation

The proposed method has been tested on all image resolutions from the e-optha database and compared to the manual measurement, based on the distance between the center of the Optic Disc and the Fovea, and the classical measurement, which uses the image width directly (Tab. 3.1). There are 8 different image sizes, and the first line (in red) is taken as the reference. The coefficient is obtained by dividing the reference by the current measurement. For each resolution, 5 images are used to get a mean value. The results show an excellent agreement between the manually measured and automatically calculated scale coefficients. In contrast, the image width is proved to be a bad reference. The proposed method shows a high accuracy and stability for estimating the scale coefficients. The generalization to other capture angles is straightforward.

3.4 Conclusions and perspectives

The issue of spatial calibration of fundus images has never been deeply studied in the literature. However, in clinical conditions, the algorithms need to deal with images with different resolutions, and sometimes the images' metadata are lost during the transmission. In this chapter, a new spatial calibration method is proposed and proved to be accurate. The width of FOV is used as invariant to estimate the size of other structures and lesions. Thus, the re-sampling of images is not necessary. However, if

the images are resized to an arbitrary size, the following algorithms will remain unchanged, because the same spatial calibration based on the new width of FOV gives new coefficients automatically.

Note that the proposed method can be generalized to any capture angle, as long as the capture angle is known, through some simple trigonometric operations (see Fig. 3.6). Let us suppose that two capture angles α and β are used, and let d_{FOV_α} and d_{FOV_β} be the corresponding FOV widths. Then we have the relation:

$$d_{FOV_\beta} = \frac{\sin(\alpha)}{\sin(\beta)} d_{FOV_\alpha} \quad (3.4)$$

This generalization has yet to be validated on clinical databases.




Figure 3.6: Illustration of the conversion of spatial calibration coefficients between the images with different capture angles.

4.1 Introduction

Exudates are one of the main clinical signs of the presence of Diabetic Retinopathy. Exudates appear as white/yellow “soft” structures in color retinal images. Their size is variable: they can be as small as a microaneurysm (i.e. a few pixels on a digital image), or as large as the optic disc. Moreover, they also show a large variability in shape and contrast. Fig.4.1 shows an example of exudates on a color retinal image and the manual annotation provided by an ophthalmologist.




Figure 4.1: Color retinal image, zoom and manual annotation produced by an ophthalmologist.

After presenting the main difficulties for exudates detection and a review of the state of the art methods, we describe in the next section the proposed exudate detection method, which is robust with respect to image variability, and artifacts. In order to achieve these goals, it automatically estimates the size parameters using the spatial calibration introduced in Chapter 3, and uses operators, including the ultimate opening, which are contrast invariant. Moreover, specific modules have been developed

CHAPTER 4. EXUDATES DETECTION

to detect artifacts, such as reflections. Finally, dozens of features are computed and a random forest is used to extract exudates among all the detected candidates. It has been validated on several databases, including e-ophtha EX.

Concerning e-ophtha, and clinical databases in general, the main problems are caused by the heterogeneity of the database, in terms of image size, color, contrast, etc. For instance, images from the e-ophtha EX database show four different resolutions, from 1440×960 pixels, to 2544×1696 pixels. Most healthy images have been chosen because they contained structures that could mislead exudate detection methods. For example, several images contain reflections along the main vessels, typical on young patients (see for example Fig.4.2(a)), and artifacts, as shown in Fig.4.2(b)).




Figure 4.2: Examples of fundus images in e-ophtha EX.

Exudates size, contrast and shape variability constitute another difficulty. Moreover, regions between two dark structures can cause false detections. All these problems will be treated in the proposed method.

4.2 State of the art

The automatic detection of exudates in color fundus images remains an open question. In the literature, recent approaches usually start with a normalization of the image and the removal of the Optic Disc (OD). Then, a first detection method provides a set of candidates, i.e. structures which are similar to exudates. Finally, a selection or a classification procedure is applied, based on features computed on each candidate, in order to keep only exudates. Machine learning methods are typically used for this last step.

Region growing methods are frequently used to obtain the set of candidates (see for example Sinthanayothin et al. [94]). Usher et al. [105] proposed a method combining recursive region growing and an adaptive intensity thresholding to get the candidates. Sánchez et al. [86] use a mixture model to fit the histogram and estimate the mean and variance of a normal distribution. The estimated parameters are used to perform a dynamic thresholding. After computing a rough detection of exudates, Walter et al. [108] use a morphological reconstruction to obtain a more precise localization of the exudate boundaries. The same method has been used by Sopharak et al. [96] and Harangi et al. [44]. We have also adopted it to refine the result of our exudate candidates segmentation method.

The risk of each candidate being an exudate can be modelled as a probability map. Walter et al. [108] use the local variance to estimate this probability, an easy to implement and computationally efficient method. Giancardo et al. [40] and Sánchez et al. [86] use the Kirsch operator to estimate it. They show that its value along the edges of each candidate is a good feature to predict the risk. Giancardo et al. [40] also use stationary Haar wavelets.

More complex machine learning methods can be used for candidates selection. Different sorts of features have been proposed in the literature. Pixel-wise features have been the first to be used. Niemeijer et al. [72], Sopharak et al. [97] and Harangi et al. [44] use pixelwise features like standard deviation, original intensity, hue channel, Sobel edge detector response, or differences of gaussians. In these cases, the authors have used fuzzy C-means and naive Bayes as classifiers. One disadvantage of this approach is that it implies a large computational load. In order to improve on this aspect, one has either to reduce the number of pixels to be classified, or to reduce the number of features. Gardner et al. [38] proposed another approach: the image is divided into small squares, which are assigned to a class (vessel, exudates, background, etc.) by a neural network.

Features based on connected components as candidates have also attracted some interest. Connected component features can be directly derived from pixel-wise features, but supplementary features can also be used, as area, length or perimeter. This approach was used by Fleming et al. [28] and Sánchez et al. [87]. Sánchez et al. [88] were the first to use contextual features in this domain. These features take into account the environment of a given candidate. Their system has two classification steps. First, local features are extracted and a linear discriminant classifier is chosen to perform the classification and obtain the candidates. Second, contextual features are calculated, such as the distance between the current candidate to the closest candidate, or the distance to a given anatomical element. The same classifier is used to perform a second classification. Results show that the performance is largely improved by using contextual information. Furthermore, Giancardo et al. [41] proposed an image level classification method. The system classifies an image into one of two classes: “healthy” or “presence of diabetic macular oedema”. Based on their previous work, Kirsch’s edge filter is used to get the candidates. Meanwhile, wavelet analysis and different color decompositions are applied to the original image. Then, within the binary mask of candidates, the global mean and standard deviation of the results of the filtered images are taken as features for the entire image. Each image has one feature vector, which is passed to a support vector machine classifier. The resulting method

is computationally efficient. Furthermore it does not need a lesion-level ground truth, but only an image-level ground truth, which is enough for screening applications.

The presented methods have been evaluated on publicly available databases. These databases are relatively homogeneous. For instance they contain mainly good quality images, taken with the same retinograph. For example, when dealing with images of different sizes, these methods deal with this variability by sub-sampling the largest images to get back to the same standard input image size, thus losing information. These characteristics make existing methods inappropriate for a telemedical mass screening application, where images come from different sources, and are therefore very heterogeneous.

4.3 Pre-processing

Given the complex problems introduced by the clinical database we are dealing with, preprocessing will not only deal with noise removal, but also with more complex problems, such as the detection of reflection zones. We proceed in two steps in order to remove spurious structures. First, we remove all dark structures, including vessels and dark lesions. Second, we get rid of bright artifacts. A clean image will be obtained after this process, simplifying exudate candidates extraction.

Two original ideas constitute the main novelties of this section:

- Bright structures, including reflections, are removed by an adaptive template;
- Bright regions along the borders of the field of view are segmented using the blue channel of the image.

4.3.1 Dark structures removal

Dark structures, including vessels and dark lesions, induce local intensity variations, which can mislead exudate candidates detection methods based on local contrast. We propose a morphological inpainting to remove dark elements. Let I_{orig} be the original image, and γ_B^n and φ_B^n respectively the morphological opening and closing of size n , with structuring element B . In this work, a hexagonal structuring element is used, as it is more isotropic than a square structuring element, while remaining computationally efficient. In order to compute the inpainted image I_{inp} , we proceed in two steps:

$$I^* = \gamma_B^{n+1}(\varphi_B^n(I_{orig})) \quad (4.1)$$

$$I_{inp} = I_{orig} \vee I^* \quad (4.2)$$

where \vee denotes the supremum operator. Fig.4.3 shows an example of inpainted image. Vessels, as well as dark and small structures, are removed. Moreover, given that we want to get rid of dark structures up to the size of the largest vessels, we choose $n = d_{vessel}$ (the largest vessels diameter, see Section 3.2).

4.3.2 Removal of bright structures

Besides the OD, there are other bright structures in fundus images which can mislead exudate detection methods. Fig.4.4 shows three kinds of common bright structures. Arrow a shows optic nerve fibers, which come out of the OD and are mainly visible along the main vessels. Arrow b points at reflections in the middle of the vessels. Arrow c shows reflections which are especially common on young patients,




Figure 4.3: Dark structures removal based on morphological inpainting. a) Original image, b) Inpainted image, c) and d) Some details.

given the extreme transparency of the eye lens and vitreous humour. This kind of reflections is usually close to the vessels. They will be taken into account in the classification section thanks to a specific feature: the distance of a candidate exudate to the nearest vessel.

We propose to use an adaptive template to remove most bright structures. See for example the simple case pictured in Fig.4.5, where we see a profile (I_{orig}) containing an exudate c and some bright structures a and b . If we can generate a template (T) like the red profile in Fig.4.5, which has higher values in the reflection regions, plus a restricted region where all maxima are removed, the reflections can be easily removed by combining a subtraction with a morphological reconstruction:

$$I_{rec} = Rec_{I_{orig}}(I_{orig} - T) \quad (4.3)$$

where $Rec_X(Y)$ means reconstruction under X by taking Y as marker. The result I_{rec} is the third profile in Fig. 4.5.




Figure 4.4: Three kinds of bright structures: a) Optic nerve fibers, b) and c) Reflections common on young patients.

In order to recover the original shape of exudates, another reconstruction is done by:

$$I_f = Rec_{I_{orig}}(Maxima(I_{rec})) \quad (4.4)$$

where $Maxima()$ extracts the regional maxima, and sets them to maximum intensity (255). The green line (I_f) is the expected result.

For different kinds of bright structures, the templates are generated by the following methods.

Firstly, the normalized local mean value is used to remove background non-uniform illumination effects. The global mean value (within the FOV) is calculated from the original image, denoted μ . Then, a mean filter with a window size of $d_{FOV}/10$ is applied on the same image. The result is named I_{mean} . The first template is given by:

$$T_{mean} = \frac{\mu I_{mean}}{\max(I_{mean})}, \quad (4.5)$$

where $\max(I_{mean})$ is the maximum grey level value of image I_{mean} .

Secondly, template T_{vessel} , corresponding to reflections within the vessels, is simply obtained by using the segmented vessel mask (details refer to Appendix B). The value of the mask is set to μ (the same global mean value computed from the original image).

Most reflections or optic fibers are found along the main retina vessels, as shown in Fig.4.6(a). They lie inside a parabolic region passing through the OD. Therefore, inspired by work done by Tobin et al. [103], we fit a symmetric double parabolic region passing through the OD center, to the largest retina vessels. A simplified version of the method is implemented. The parabola fitting is based on choosing the best parameter from a predefined set of parameters. First, we remove from the vessel mask the branches which are less than $(2/3)d_{vessel}$ pixels wide. Then a parabola is formulated as:

$$x - x_{OD} = a(y - y_{OD})^2 \quad (4.6)$$

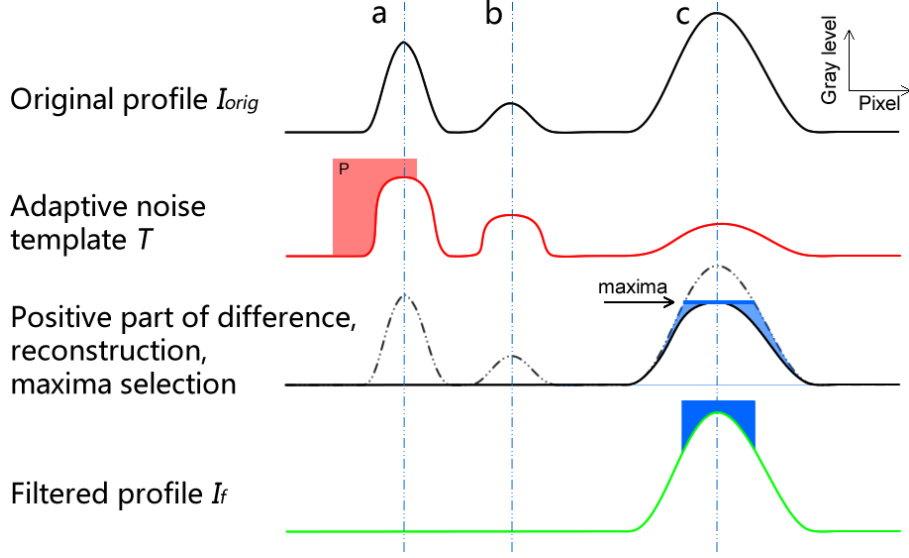


Figure 4.5: Adaptive template for the removal of bright structures - illustration on an image profile. a and b are bright structures, and c is an exudate. The bright structures can be removed by the following steps: subtract the template from the original profile and keep the positive part. Perform a reconstruction to restore the profiles of remaining structures. Maxima touching the restricted region P are removed. Another reconstruction by the maxima is used to restore the exudate.

where parameter a is selected from $\{0.02, 0.013, 0.01, 0.005, 0.003, 0.0025, 0.002\}$. x_{OD} and y_{OD} are the coordinates of the OD center. The generated parabola is dilated in the horizontal direction by a linear structuring element of width $1/15d_{FOV}$, to obtain a parabolic region. Thus, according to the parameters list, 7 parabolic regions are generated. The one containing the most filtered vessels is selected. The mask is denoted $T_{parabola}$. Its value is set to μ . Within this region, we define a “forbidden” sub-region P corresponding to the parabolic region which bends around the macula, up to a horizontal distance to the OD equal to $d_{FOV}/4$, where the value is set to the maximal authorized value of the image (here 255). This region is depicted in white in Fig. 4.6(b).

The final template is given by the supremum of the first three templates:

$$T_{final} = T_{mean} \vee T_{vessel} \vee T_{parabola}. \quad (4.7)$$

Fig. 4.6(b) illustrates the result.

This template image is used as reference to analyze the maxima of the inpainted image: All regional maxima that are lower than the template image T_{final} , or that touch region P , are erased. The final result will be referred to as the **preprocessed image**. This treatment is illustrated in Fig. 4.6. Note that all bright structures which lay in regions where reflections and optical fibers are currently found are suppressed when their value is lower than the template value at the same position.

Some images contain excessively bright regions along the FOV border (see for example Fig.4.7(a), red arrow). These bright regions can cause false detections. To deal with them, they are segmented using the blue channel during the OD detection (For details, refer to Appendix C). This mask will be used during candidates extraction.




Figure 4.6: Reflection level estimation and computation of preprocessed image. a) Original image, with overlaid parabolic region, b) Adaptive template, c) Inpainted image, d) Final result, obtained after removing maxima which touch region P , or that are lower than the template image.

4.4 Candidates extraction

In this section, we present a two-scale exudate candidates segmentation method. Large exudate candidates are obtained from the preprocessed image, thanks to a mean filter followed by a reconstruction. Small exudate candidates are directly computed on the green channel of the original image, by means of a morphological top-hat.

In the case of large exudates, firstly we compute a mask, which is the union of the FOV contour, the OD mask (modelled as a disk), and bright border regions (if any). This mask is taken as marker to perform a reconstruction under the preprocessed image (Fig. 4.8(a)). A comparison between the reconstructed image and the preprocessed image is done, and pixels where the values are equal in both images are set to zero. The result is illustrated in Fig. 4.8(c). This step keeps only relevant regions. The candidates will be extracted only from these regions. Secondly, a mean filter with window size equal to $d_{FOV}/10$ is applied on the preprocessed image, and the result is subtracted from the relevant




Figure 4.7: Bright border regions segmentation. a) Blue channel, b) Result of segmentation.

regions image; only the positive part of the result is kept. Finally, a threshold of 10 gives the rough candidates mask (Fig. 4.8(d)).

In order to get a more precise shape of the candidates - an important feature for the subsequent candidates classification - an efficient technique based on a morphological reconstruction, initially proposed by Walter et al. [108], is used. We begin by slightly dilating the obtained rough candidates mask, using a structuring element of size $d_{vessel}/2$. All pixels belonging to the resulting mask are set to zero in the original image, as shown in Fig. 4.9(a). Then, a morphological reconstruction is applied by taking the previous image as marker and the original image as mask (Fig. 4.9(b)). Finally, the difference between the reconstructed image and the original one gives the **large exudate candidates image** (Fig. 4.9(c) for an illustration).

Small exudate candidates are missed by this method. In order to detect them, we apply a morphological top-hat with a structuring element of size d_{MA} to the green channel of the original image. The resulting image contains the **small exudate candidates** (Fig. 4.9(d)). The supremum between this image and the large exudate candidates image gives the **candidates contrast image**. This grey level image will be extremely useful when computing candidate descriptors, as it contains precious information on their shape and contrast. Finally, by thresholding the candidates contrast image at level 5, and removing all connected components which contain less than 5 pixels (the size is fixed for all types of images because noise is almost resolution independent) we obtain the **exudate candidates mask**. Each connected component of this binary image is considered in the following section as an exudate candidate.

4.5 Classification and risk evaluation

The candidates set contains exudates, but also other bright structures. In order to classify the candidates (i.e. the connected components of the candidates mask), we adopt a classic strategy: we compute features on them, and then use a machine learning technique – here a random forest method (Breiman [13]) – to classify them.




Figure 4.8: Large exudate candidates extraction illustration. a) Preprocessed image, b) FOV contour, OD and bright border regions mask, c) Relevant regions, d) Rough mask of large candidates.

4.5.1 Features

Two novelties are introduced in this section:

- some features are not only computed on the whole candidate, but also on a smaller support, which corresponds to a different level of the candidates contrast image;
- new contextual and textural features are introduced.

The features are based on the following images: the candidates mask, the candidates contrast image (*contrast*), the inpainted image (*inp*), the preprocessed image (*pre*), and the green and saturation channels of the original image. A typical feature name results from appending the code corresponding to the image taken into consideration, with the code corresponding to the measure. For example, *meanContrast* corresponds to the mean value of the candidates contrast image. The final value is obtained by integrating the measure on the considered image over the exudate candidate. In some cases, the image is not specified in the feature name, either because it is implicit (like for geometric features, e.g. *perimeter*), or for simplicity.

As introduced in Section 2.2.1, we have discovered that in some cases considering different levels for the same exudate candidate brings useful information into the classification. In this case, instead




Figure 4.9: Precise shape of candidates extraction and small candidates. a) Green channel with candidates set to zero, b) Reconstruction, c) Large exudate candidates, d) Small candidates .

of a single feature, such as $varInp$, we have two derived features $varInp_1$, computed on the basic layer, and $varGreen_2$, computed on the second layer.

Intensity features

The maximum, minimum, mean and median values of each candidate (coded respectively as max , min , $mean$ and $median$) in the candidates contrast and $green$ channel images, give the first basic intensity based descriptors. For example, $maxContrast$ gives the maximum value of the candidates on the residual image. Note that $minContrast$ is not taken into account, as its value is always equal to 1 in practice. Descriptor $diffGreen$ is the difference between $maxGreen$ and $minGreen$.

The saturation channel is used to distinguish another sort of optical artifacts, which can be introduced by the camera optics. They can appear anywhere on the image, as shown in Fig. 4.10(a). They are difficult to distinguish from normal exudates in any of the RGB channels of the image (Fig. 4.10(b)), but we have found that they can be easily detected in the saturation channel (Fig. 4.10(c)).

The saturation values vary between different image sources. But for the same type of image, the reflections are darker than other bright structures in the saturation channel. Thus, we normalize the saturation channel by dividing it by its global mean value. The final descriptors $meanSat_1$ and

$meanSat_2$ are obtained by computing on layers 1 and 2 of each exudate candidate the mean value of the normalized saturation channel.




Figure 4.10: Optical artifacts. a) Color image with three artifacts and small exudates, b) green channel, c) saturation channel

Geometric features

Descriptors *area* and *perimeter* are obtained by counting the corresponding numbers of pixels. *integralContrast* is the accumulation of the intensity in the candidates contrast image. These three size dependent features are normalized by dividing them by d_{vessel}^2 . The circularity of an exudate candidate is given by:

$$circ = \frac{4A}{\pi L^2}, \quad (4.8)$$

where A is its area, and L its geodesic diameter (see Morard et al. [70]).

Textural features

Local variance is used to estimate the contrast change. It is calculated in a sliding square window of width d_{vessel} . The result is shown in Fig. 4.13(a). We compute it on the inpainted image I_{inp} , in order to avoid the variance introduced by dark structures. It is moreover computed at two levels, leading to two features $varInp_1$ and $varInp_2$.

The gradient is classically used for edge characterization. Here, we use it to characterize the roughness inside a given region. After computing the morphological gradient of the inpainted image with a hexagon of size 2 inside a given exudate candidate, we compute the value n of the flooding of the region until the minima merged into the background or the size of them exceeding d_{MA}^2 pixels. We call the resulting descriptor *swamping*, as the name of this procedure in mathematical morphology. Fig. 4.11 illustrates the method on a profile. Fig. 4.12 shows an example on the fundus image which

has several exudate candidates. Fig. 4.12(a) is the inpainted image. After computing the gradient (Fig. 4.12(c)), the regional minima are shown in Fig. 4.12(d). Note that the entire background is a large minimum in this image. Thus, with the increasing of the swamping level, the minima inside the candidates are merged into the background successively, for example, Fig. 4.12(g) shows the minima of the result of the swamping with height 10. An area selection keeps the small minima (Fig. 4.12(e) and Fig. 4.12(h)). For each candidate, the level when all minima inside it are removed, is noted as the *swamping* feature.




Figure 4.11: Illustration of computing the swamping feature. Top: original profile of an exudate; Middle: its gradient, in red, and the result of adding a constant n to the gradient. Bottom: when n is large enough, after reconstruction all regional minima are removed.

A hybrid feature: the ultimate opening

The ultimate opening is a multi-scale morphological operator introduced by Beucher [11]. It extracts the most contrasted structures and the corresponding size information from an image. More precisely, for each pixel the ultimate opening (θ) keeps the largest difference, denoted R , between consecutive openings, as well as the size of the opening corresponding to this largest difference, denoted q . The UO is defined as:

$$\theta(I) : I \xrightarrow{\theta} (R(I), q(I)) \quad (4.9)$$

$$\begin{aligned} R(I) &= \vee (r_\lambda(I)), \forall \lambda \geq 1 \\ \text{where } r_\lambda(I) &= \gamma_\lambda(I) - \gamma_{\lambda+1}(I) \end{aligned} \quad (4.10)$$

$$q(I) = \begin{cases} \vee \{\lambda + 1 \mid \lambda \geq 1, r_\lambda(I) = R(I)\} & R(I) > 0 \\ 0 & R(I) = 0 \end{cases} \quad (4.11)$$

where $(\gamma_\lambda)_{\lambda \in \{1, \dots, N\}}$ is an increasing family of openings. In our case, only the residue R of the preprocessed image is considered. We apply the ultimate opening to the inpainted image (Fig. 4.13(b) for an illustration). For each candidate, the mean value of $R(I_{Inp})$ within the exudate candidate is computed. Two derived features are obtained: $uoInp_1$ and $uoInp_2$.




Figure 4.12: Illustration of computing the swamping feature on fundus image. a) Inpainted image, b) Exudate candidates, c) - h) Three columns correspond to morphological gradient, regional minima of the gradient and area selection of the regional minima. The first row is computed on the inpainted image. The second row is computed on the swamped image with the height of 10.)

Contextual features

Reflections and optic fibers lay often beside the vessels, as shown in Fig. 4.14. This information provides contextual information, as introduced in Sánchez et al. [88]. $distCenter$ is the distance between the barycenter of the candidate to the nearest vessel. $distMin$ is the minimum distance of the candidate to the nearest vessel. They are normalized by dividing them by d_{vessel} .

Exudates sometimes appear in a bunch within a small region. Thus, if more than one high risk candidate appears, the risk associated to this region should be raised. Sánchez et al. [88] did a first classification by using local features to assign a class (lesion, background etc.) to each candidate. Then, features as the number of candidates within a neighbour region and the distance to the nearest candidates are extracted. In order to simplify this process and keep the original information as much as possible, we propose using the area and the number of regional maxima in the neighbor region of




Figure 4.13: a) local variance, b) ultimate opening computed on the preprocessed image




Figure 4.14: Reflections beside vessels

a candidate, which provide similar contextual information. First, we define the neighbor region by a window, which is based on the bounding box of the candidate in the Cartesian system, while each direction is expanded by $2d_{vessel}$. Within the window, the original green channel is thresholded at $meanGreen$ and at $minGreen$. As shown in Fig. 4.15, the red part is the subset by thresholding at $meanGreen$, which has two maxima ($nbMaxH$). The area ($areaH$) is defined by the total number of pixels inside the window after thresholding (excluding the candidate itself). $nbMaxL$ and $areaL$ are calculated by the same method at the lower threshold level (the blue plus red part in Fig. 4.15).

Final set of features

Finally, twenty eight features are computed on each candidate. They are grouped in Table. 5.2.




Figure 4.15: Neighbor structures. An exudate is in the middle of the image with four surrounding structures. Two thresholds are used to select the neighbour structures. The number of the maxima and area are noted.

Intensity:	<i>maxContrast, meanContrast, medianContrast, minGreen, maxGreen, meanGreen, medianGreen, diffGreen, meanSat₁, meanSat₂</i>
Geometric:	<i>area₁, area₂, volume₁, volume₂, perimeter, circularity</i>
Textural:	<i>varInp₁, varInp₂, swamping</i>
Hybrid features:	<i>uoInp₁, uoInp₂</i>
Contextual:	<i>distCenter, distMin, nbMaxH, nbMaxL, areaH, areaL</i>

Table 4.1: Feature list

4.5.2 Classification

Random Forest (Breiman [13]) is used to perform the classification of the exudate candidates. The number of trees is set to 500. It can evaluate the importance of features during the generation of the forest. Because of the probabilistic nature of random forests, each run gives a slightly different result. Fig. 4.16 shows the result which is obtained by taking the mean value of 5 runs. The sum of all the values is equal to 1. We can observe that:

- Contextual features play an important role in the classification. Indeed, among the first five features, three are contextual. This result confirms those of Sánchez et al. [88]. Note also that our new contextual feature *nbMaxH*, appears in the top 5.
- The new hybrid features *uoInp1* and *uoInp2* based on the ultimate opening appear also in the top 5.
- The new features *meanSat₁* and *meanSat₂* based on the image saturation obtain good scores. This good result is probably linked to their discriminant power with respect to reflections.
- The raw intensity features achieve a middle rank. Even if the OD has been removed from the candidates set, other bright structures witch are not exudates remain in the image. This explains why the intensity features have a low discriminative power compared to others.

- Most geometric features (*area*, *perimeter* etc.) are not very interesting for exudates classification. This result is not unexpected, as exudate size and perimeter show great variability. But *circularity* is an exception, because most reflections have an elongated shape.




Figure 4.16: Features importance

The Random Forest algorithm gives a probability of being an exudate to each exudate candidate. Fig. 4.17, shows such a probability map, which can be compared with the manual annotation.

4.6 Results

In this section, we evaluate the proposed method on the new database, e-optha EX, as well on three other publicly available databases: Messidor, DiaRetDB1_v2 and HEI-MED. The evaluation of the results can be done at two different levels. Firstly, when an accurate delineation of the exudates




Figure 4.17: Risk function (exudates). a) Original image, b) Manual annotations, given by the e-ophtha database, c) Probability map

is available, the evaluation can be done at the individual exudate level. This kind of evaluation is pertinent for applications where a precise quantification of the lesions is important. For screening applications, it is often enough to evaluate the results at the image level (i.e. presence of absence of lesions).

4.6.1 Materials

Besides e-ophtha EX, there are several publicly available databases for the evaluation of exudate detection methods. Messidor (Dupas et al. [24]) is a large database containing 1200 images. It provides a DR diagnostic for each image, but does not contain exudate contours. The DIARETDB1 database, introduced by Kauppi et al. [55], contains 89 images, and provides rough exudates annotations. However, these annotations are not precise enough to evaluate a segmentation algorithm at the lesion level. HEI-MED (Giancardo et al. [41]) is a database dedicated to train and test image processing algorithms for the detection of exudates and diabetic macular oedema. It contains 169 eye fundus images. Each image of the database was manually segmented by an expert. But it shows the same problem as DIARETDB1: the segmentation is not precise enough for a lesion level evaluation.

4.6.2 Lesion level validation

To the best of our knowledge, this is the first time that an exudate segmentation method can be evaluated at the lesion level on a publicly available database, containing precise lesions annotations. This is possible thanks to the new e-ophtha EX database.

The evaluation can be classically done by counting the number of pixels which are correctly classified. However, as other authors, we considered this approach inappropriate for exudate segmentation evaluation. Indeed, imagine the situation shown in Fig. 4.18. In the middle, there is a detected exudate (in blue) and the corresponding ground truth (in red). In practice, most people would consider that this exudate was correctly detected, even if the contours do not match perfectly. If we only count the pixels belonging to the intersection as true positives, we would get half blue pixels as false positives and half red pixels as false negatives. Moreover, this kind of measure would tend to under-estimate

errors on small connected components. This is the reason why authors such as [40] and [88] have resorted to connected component level validation: a connected component candidate is considered as True Positive (TP) if, and only if, it touches the ground-truth. Thus the connected components of the segmentation and of the ground truth can be classified as True Positives (TP), False Positives (FP) and False Negatives (FN). However, attributing the same weight to a large exudate and a small exudate seems inappropriate. Moreover, with this approach a single very large detection mask would produce excellent results as long as it covers the whole ground truth set. Therefore, a hybrid validation method is proposed below, where a minimal overlap ratio between ground truth and candidates is required.




Figure 4.18: Illustration of the definition of True Positive and False Positive pixels. In blue: detected candidates; in red: ground-truth

The main problem is therefore how to appropriately fusion the overlapped detected candidates and ground-truth sets. To proceed with this evaluation, we take as starting point evaluation methods used for document analysis introduced by Wolf and Jolion [110]. The set of candidates is $\{D_1, D_2, \dots, D_N\}$, where D_i corresponds to a connected component. Similarly, the set of ground truth exudates is $\{G_1, G_2, \dots, G_M\}$. The masks of candidates and exudates are therefore respectively

$$D = \bigcup_{1 \leq i \leq N} D_i \quad (4.12)$$

and

$$G = \bigcup_{1 \leq j \leq M} G_j. \quad (4.13)$$

A pixel is considered as a True Positive, if it belongs to any of the following sets:

- $D \cap G$;
- D_i such that $\frac{|D_i \cap G|}{|D_i|} > \sigma$;
- G_j such that $\frac{|G_j \cap D|}{|G_j|} > \sigma$;

where $|\cdot|$ is the cardinal of a set, and σ is a parameter belonging to $[0, 1]$. If it is equal to 0, any D_i (resp. G_j) touching G (resp. D) will be considered as True Positive.

A pixel will be considered as a False Positive, if it belongs to any of the following sets:

- D_i such that $D_i \cap G = \emptyset$;

Threshold	TP pixels	FN pixels	FP pixels	Sensitivity	PPV
0	293,466	60,331	305,279	83%	49%
0.4	261,771	92,026	99,569	74%	72%
0.8	103,311	250,486	9,396	30%	92%

Table 4.2: Details of pixel level validation.

- $D_i \cap \overline{G}$ such that $\frac{|D_i \cap \overline{G}|}{|D_i|} \leq \sigma$.

A pixel will be considered as a False Negative, if it belongs to any of the following sets:

- G_j such that $G_j \cap D = \emptyset$;
- $G_j \cap \overline{D}$ such that $\frac{|G_j \cap \overline{D}|}{|G_j|} \leq \sigma$.

All other pixels are considered as True Negatives (TN).

Given that these four classes are, in our case, clearly unbalanced, as TP, FN and FP are in practice negligible with respect to TN, computing the specificity, i.e. $\frac{TN}{FP+TN}$, and therefore a ROC (Receiver operating characteristic) curve, does not seem appropriate.

We will finally compute the sensitivity S of the detection on the one hand:

$$S = \frac{TP}{TP + FN}, \tag{4.14}$$

and the positive prediction value (PPV) on the other hand:

$$PPV = \frac{TP}{TP + FP}. \tag{4.15}$$

Parameter σ has been set to 0.2. We chose a small value in order to correctly take into account situations such as the one depicted in Fig. 4.18. We did not want to use $\sigma = 0$, which would correspond to the approach proposed by [40] and [88], with the drawbacks described above. Note however that the exact value of σ is not critical, as our main objective is to compare different methods – or to optimize a given method. Fig. 4.20(b) shows an evaluation result on an image. Note the FNs in the image bottom. Only a few points are detected as candidates; not enough to satisfy the criterion associated to $\sigma = 0.2$. Thus, the other pixels in the connected component of the ground-truth are regarded as FNs.

The test is done on the e-ophta EX database with leave-one-out cross-validation. Candidates are grouped by image. The candidates from the test image are left out, while the rest are used to train a model. As we increase the threshold on the probability given by the classification to obtain the detection mask, the sensitivity decreases, and the PPV tends to increase (note however that PPV is not an increasing function of the threshold). Fig. 4.19 sums up the result.

Table. 4.2 gives some values, corresponding to different probability thresholds. We can see that a number of pixels are not detected even with the lowest threshold. These are mainly small low contrasted exudates. Fig. 4.20(d) shows such a case, containing 2400 TP pixels and 1037 FN pixels. Some FPs are due to the presence of other bright lesions, like cotton wool spots and drusens, which are often considered as exudates. This is not a major problem for a screening application.




Figure 4.19: Sensitivity and precision (positive predictive value) for pixel level evaluation.




Figure 4.20: Example of pixel level validation. a) and c) Original images, b) and d) Result of pixel level validation with $\sigma = 0.2$, and minimum probability threshold.

Before proceeding to the next section, we would like to insist upon the fact that the exudate segmentation evaluation method presented in this section aims at evaluating and comparing exudate detection algorithms. For a clinical evaluation, we believe that image level evaluation, introduced below, is more appropriate.

4.6.3 Image level validation

The result of our exudate detection method is a set of connected components, each of them accompanied by a probability. In order to compute an image level probability, i.e. a probability that the image contains at least one exudate, we take the maximum of all individual probabilities.

In order to evaluate our algorithm on the e-optha EX database a leave-one-out cross-validation is adopted. The resulting ROC curve is shown in Fig. 4.21. The corresponding AUC (Area Under the Curve) is 0.95. For example, with a specificity of 89%, we obtain sensitivity of 96%. This is a very satisfactory result from a diabetic retinopathy screening point of view.

We have applied to the same database the method proposed by Giancardo et al. [40], one of the best state-of-the-art methods for exudates segmentation, which was developed for the HEI-MED database. The resulting ROC curve is given on the same diagram. The AUC is 0.87. Our method gives better results on e-optha, which is natural as it has been designed for this kind of data. For example, to

achieve the same 96% sensitivity with the method by Giancardo et al. [40], specificity would drop to 63%.

In order to evaluate the robustness of our method, we have applied it to the other three publicly available databases, DiaRetDB1_v2, HEI-MED and Messidor, after learning on the whole e-ophtha EX database. Again, note that the image processing (or the candidates extraction) part is automated, for example, the parameterizations of the filters and the thresholdings. We do not change anything while performing the test on the other databases, except for DiaRetDB1_v2, because their images are acquired with a field of view angle 50° . We adjust the coefficients of spatial calibration. The resulting ROC curves are given on Fig. 4.22. The corresponding AUC values show that the performance remains similar. This result shows that our method is robust with respect to changes in the acquisition conditions.

In Giancardo et al. [41], the authors proposed a method to predict the probability of presence of an exudate per image. The method was tested on the three previously cited databases. Table 4.3 sums up the results, where learning and testing have been done on the same database, using leave-one-out cross-validation. In spite of the fact that our method has been optimized for the e-ophtha EX database, we can see that it obtains better or similar results than the method by Giancardo et al. [41]. Moreover, the performance of this method decreases when learning is done in one database, and testing on another database. For example, according to Giancardo et al. [41], a model trained on Messidor and tested on HEI-MED gives an AUC 0.82. This shows that the proposed method, developed for a heterogeneous database, tends to be more robust than existing methods when directly applied to other databases.




Figure 4.21: ROC curves and AUC values of proposed method with cross validation and method from Giancardo et al. [40] on e-ophtha EX.




Figure 4.22: ROC curves and AUC values of proposed method with cross validation on e-ophtha EX and validation on other three public database with the model trained on e-ophtha EX.

	Proposed method	Giancardo et al. [41]
DiaRetDB1_v2	0.95	0.93
Messidor	0.93	0.90
HEI-MED	0.94	0.94

Table 4.3: Comparison of AUC of three public database. The proposed method, optimized for e-ophtha EX, has been directly applied to the other databases (except for a small adaptation for DiaRetDB1_v2, see Section 4.6.3). The method by Giancardo et al. [41] has been individually optimized for each database

4.7 Conclusion

In this chapter, a new method for exudates segmentation on color fundus images has been introduced. The method is able to deal with images showing large variability in terms of quality, definition and presence of artifacts. It is, to the limit of the authors knowledge, the first method to be able to successfully process images containing reflections, which are frequently found on young patients. In order to reach this objective, the proposed method combines a precise candidates segmentation step, with a classification step, where new efficient features are introduced.

The proposed method is tested on the new data base, e-ophtha EX, which contains on the one hand images where the exudates have been accurately contoured and, on the other hand, healthy images, without exudates. The proposed method has been validated on this database, and has been compared with state-of-the-art methods on other publicly available databases. The results show that the new method not only performs better than the previous ones, but that it is very robust, as the results remain good on all tested databases, without changing any parameters.

We have invested a great effort into reducing the number of parameters as much as possible. The fact that the results are competitive not only in the e-ophtha EX database, which was the primary

CHAPTER 4. EXUDATES DETECTION

aim of this study, but also on other publicly available databases, without any specific learning or parameter tuning, shows that this goal has been reached. However, there are still some parameters left, mostly corresponding to thresholds. One of our current goals is to use a grey level calibration of the images, based on the same philosophy as our spatial calibration, in order to automatically adapt these parameters to the image content.

5.1 Introduction

Microaneurysms are likely to be the first lesions present at the earliest stage of Diabetic Retinopathy. Their detection is important for the diagnosis of disease in its early stage. However, because of their small size and sometimes low contrast, the detection can be easily disturbed by noise or fragments of structures like vessels. This makes their automatic detection difficult. In this chapter, we are going to first talk about the characteristics of microaneurysms supported by statistics from a rich database. In the following section, after a brief review of the state of the arts, we will propose a new method to detect microaneurysm candidates. New descriptors are proposed to extract features and Random Forest classification is used to assign risk values to each candidate. In the end, we will evaluate the method and conclude.

5.2 Microaneurysm characteristics

Microaneurysms are dilations of the venous end of retinal capillaries. Their sizes vary mostly between 10 and 100 μm , and never exceed 125 μm . In color fundus images, microaneurysms appear as isolated dark red dots (Fig.5.1). In order to get overall characteristics of microaneurysms, we did an analysis on the e-optha MA database (details are introduced in Chapter 1, Section 1.3). 148 images which contain 1305 microaneurysms and other red dot lesions (small hemorrhages) are used to perform the analysis. The main idea is to reconstruct all microaneurysms from the expert's annotations, to model the general aspect of microaneurysm, from an image processing point of view. We recall that these annotations correspond to a single point, which should belong to the MA. The statistical analysis will reveal the distributions of their contrast and size.

5.2.1 Restoration of microaneurysms

The microaneurysm contours are obtained from the specialist's markers thanks to the following steps:




Figure 5.1: A microaneurysm in color eye fundus image and its green channel.

- The proposed pre-processing method based on alternating Sequential Filter (see Chapter 2) is applied to the original green channel. The resulting image is called in this chapter the **pre-processed contrast image**. The result is illustrated in Fig. 5.2(b). All dark structures and lesions are kept. After a thresholding and removal of large objects, we get all potential candidates in Fig. 5.2(c). In the end a reconstruction from the expert's annotations restores all marked microaneurysms (Fig. 5.2(e)).
- Certain microaneurysms could be eliminated during this process because of the following reasons. First, two annotations can be very close and touch the same maxima in the image. Second, a microaneurysm may touch or intersect a large dark structure. Finally, 1294 microaneurysms out of 1305 are restored. It is enough for a statistical analysis.

5.2.2 Contrast

Contrast is the most basic but important feature to characterize a microaneurysm. The contrast of each microaneurysm is calculated on the pre-processed contrast image as shown in Fig. 5.3(a). The mean contrast and max contrast are computed. The mean contrast is calculated on the middle layer (see Fig. 5.3(b)). Fig. 5.4 shows their distribution in the e-optha MA database. The contrast of more than 90% of microaneurysms belong to the interval 5 to 15. Note that there are 14 microaneurysms whose contrast is lower than 5 which makes them almost impossible to detect.

Concerning the maximum contrast, the histogram is shifted to the right. There are still 7 microaneurysms whose maximum gray level is lower than 5.

5.2.3 Size

When we introduced spatial calibration (Chapter 3), we defined a maximum size of microaneurysm (d_{MA}) from the size of the FOV (d_{FOV}). Among 148 images of our database, we have 4 different resolutions. Therefore we have to perform a spatial calibration before extracting size information.




Figure 5.2: Reconstruction of MA. a) Original green channel, b) Negative residues of ASF (also regarded as contrast image), c) Binary mask of all candidates, d) Markers from experts, e) Restored microaneurysms.

Area of both bottom and middle layer are normalized by dividing the theoretical maximum area (d_{MA}^2) given by the spatial calibration. Thus the result is normalized to a value between 0 and 1 (For several extreme large microaneurysms, the values are limited to 1). The histogram of the entire set is plot in Fig. 5.5(a). Notice that 21 microaneurysms are under 0.1 (for example, for the resolution of 1440×960 , we have 9 microaneurysms with area less than 10 pixels) which makes them difficult to detect based on an area criteria.

A 2D model of microaneurysm profile could be created from extracted information. Indeed, a larger sampling of the layers gives a better model. Although we only extracted the area of the middle and the bottom layer, it's enough to give a preliminary model. The ratio between middle and bottom layer is calculated for each microaneurysm. Its histogram is shown in Fig. 5.5(b) and the median value of the ratio is 0.3. If we normalize the bottom layer area to 1, the middle layer would have an area of 0.3 and a width of 0.55 in the profile. First we plot the original profile of a microaneurysm (Fig. 5.6(a)). Both size and contrast are normalized to 1. The red crosses delimit the theoretical border of the middle layer. In this example, the microaneurysm is a little wider than the average. If we take the top and the middle points into account, two simple models can be given by linear and Gaussian approaches




Figure 5.3: Details of a microaneurysm. a) Microaneurysm in pre-processed image, b) Topographical view of a MA.

(Fig. 5.6(b) and Fig. 5.6(b)). With the bottom width limited from -0.5 to 0.5 , a value of 0.33 of the standard deviation of the Gaussian gives a good approximation. Both models don't reach zero at the border. Comparing them, with the limited information, we cannot say which one is better. But the Gaussian model is smoother and easier to implement as a template match filter. In order to create a better model, more layers should be extracted and a more complex mixture model should be used. In the following section, this kind of information is extracted by a multi-threshold decomposition. We won't use an explicit model, but a machine learning to select microaneurysms from the candidates.

This part analyzes two general characteristics of microaneurysm: contrast and size. Two simple models are created. A statistical study on all the microaneurysms in the e-optha MA database provides important features which will be used in the following section to build a microaneurysms detection method.

5.3 Microaneurysms detection method

5.3.1 State of the art

The first work in this domain was done by Laÿ [60] who proposed microaneurysm detection and vessel segmentation algorithms based on mathematical morphology. Microaneurysms are detected by a Top-Hat and selected by size and contrast criteria. Vessels are segmented by a rotating linear structuring element. Though the application was on fluorescein angiograms and under very limited conditions (without modern computers), the developed algorithms made an important contribution to this domain. Spencer et al. [98] proposed a complete sequence of methods to detect microaneurysms in angiograms. After illumination-uniformity correction and shade-correction, a similar Top-Hat by a rotating linear structuring element was used to remove vessels. Microaneurysm candidates were extracted by a matched filtering. Then a region growing algorithm was used to get original size and contrast information. A number of features, like aspect-ratio and energy, were calculated on each candidate. A selection was done by logically combining the criteria. Cree et al. [19] brought an improvement by redesigning the region growing part and the classification stage. Frame et al. [34]




Figure 5.4: Contrast distribution of microaneurysms.

improved previous methods by applying different machine learning methods. Yang et al. [112], Streeter and Cree [101], Niemeijer et al. [71] and Fleming et al. [30] inherited the method from Spencer et al. [98] and brought a few changes to remove false positives. Niemeijer et al. [71] combined mathematical morphology based operators and a preliminary pixel based classification method to extract candidates. New features were introduced for classification, such as color and contextual features. Fleming et al. [30] used the morphological watershed to remove vessels and bright structures, thus to extract regions of interest, while extracting features. Walter et al. [109] used a closing by diameter to get candidates, and Ravishankar et al. [83] used a morphological fill holes. Zhang et al. [117] used a method based on a new morphological operator: the Ultimate Opening.

Besides mathematical morphology, matched filtering is another popular method, because of the fact that microaneurysms may be modelled by Gaussians. Marino et al. [64] and Zhang et al. [114] used a set of Gaussian kernel correlation filters. Kande et al. [54] adopted the same approach combined with an entropy-based thresholding to get MA candidates. Bhalerao et al. [12] used a Laplacian of Gaussians and shape filtering by a circular-symmetry operator on an orientation map. Giancardo et al. [39] proposed to use a Radon Cliff operator directly on the original image, which is able to detect Gaussian-like circular structures regardless of their size or contrast. Microaneurysms are detected by using the two outputs of the Radon Cliff operator. There are other similar methods like those by Abdelazeem [3] who use circular Hough transform, and by Mizutani et al. [69] who use double-ring filter. Lazar and Hajdu [61] did a profile analysis to extract MA candidates based on their diameter and circularity.

There are also a number of algorithms with various technologies. Sinthanayothin et al. [94] and Usher et al. [105] used the Moat operator, a type of high-pass filter. Then recursive region growing and adaptive intensity thresholding are used to detect the candidates. Quellec et al. [78] proposed a method based on wavelets. A model of MA is designed in the sub-band of the image of wavelet transform by using a 2D Gaussian function. Sánchez et al. [89] used a mixture model-based clustering for candidate extraction and logistic regression for classification. Antal and Hajdu [7][8] proposed an ensemble-




Figure 5.5: Microaneurysm size distribution. a) Area after spatial calibration, normalized to a value between 0 and 1 (1 stands for the largest size of microaneurysms), b) Ratio between the area of middle and bottom layers.

based method. A set of preprocessing and microaneurysm detection methods are selected from the literature. The performances with different combinations are tested in six microaneurysm categories (based on visibility: subtle, regular and obvious; based on spatial location: near vessel, in macula and periphery). In order to combine all the operators, two approaches are proposed: search-based and adaptive weighting.

Most of the algorithms proposed in the literature have been developed either for small databases, or large homogeneous databases. Therefore, they are not suitable for the e-ophtha MA database, which contains heterogeneous images. We propose below a new microaneurysm detection method, which has been designed to be robust with respect to image quality degradation, and resolution changes.

5.3.2 Pre-processing

Contrast tends to diminish towards the edge of the images and intensity varies between images. The illumination within an image and between images are significant different. In Chapter 2, we have discussed in detail several pre-processing methods, and proposed a method based on the residue of alternating sequential filter (ASF). We use this method as the pre-processing method. However, in the literature, there are some complex and microaneurysms-oriented pre-processing methods, for example, vessel removal and noise suppression. However, they are not considered in this section. Vessel segmentation method is presented in Appendix B while details about distinguishing between microaneurysms and vessels will be discussed in the following section. Noise suppression should be treated with extreme caution, as microaneurysms can be easily considered as noise. We have tried median and Gaussian filters to remove small fluctuations, or an inpainting technique to remove bright structures. But side effect include the loss of small low contrast microaneurysms or of the ones close to vessels. In order to conserve as much as possible the original information, here we do not apply any filter which would change the image content.

One disadvantage of the proposed pre-processing method is the presence of small dark regions in the middle of high values in the residual images, which may cause a false detection. An example is




Figure 5.6: Microaneurysm models. a) Original profile, b) Linear model, c) Gaussian model.

given in Fig. 5.7. We will not treat this problem here. Instead, the classification step (Section 5.3.6) should avoid these false detections.

5.3.3 Microaneurysm candidates detection

From the above sections, we can list the main challenges of microaneurysm detection:

- Their contrast and size varies, but they are small in general. They are similar to vessel fragments, and difficult to distinguish from noise.
- Ophthalmologists are sometimes not sure that a red dot is a microaneurysm or another structure or lesion (such as dot hemorrhages).

In this section, we propose a general method to detect microaneurysm candidates. The object is to keep as much as possible the true microaneurysms. We are not looking for the best sensitivity -




Figure 5.7: A dark region among exudates leads to a rise in the pre-processed image. a) Original image, b) Result of pre-processing.

specificity balance. We leave this problem to the classification section. The detection starts from the pre-processing contrast image (see Section 5.2.1), shown in Fig. 5.9(b). We denote this image I_{pre} .




Figure 5.8: Maxima over a high platform removed by a reconstruction. Left side: microaneurysm with one maximum; Right side: a large structure (vessel or hemorrhage) with two maxima. After reconstruction, the two maxima in the right side become one large maximum which exceeds the size criteria.

Large structures removal. Each microaneurysm in I_{pre} contains at least one maximum. But large dark structures can have a lot of similar maxima, as shown in Fig. 5.8. A microaneurysm with one maximum is in the left. A high platform with two maxima is in the right. At first, I_{pre} is divided by 2. A reconstruction is done under I_{pre} (blue region in Fig. 5.8 and Fig. 5.9(c) for real condition). If the maxima over the platform are not too highly contrasted, as in this example, they will be completely razed. Then, for the left structure, we obtain a middle size maximum (compared to the maximum size of microaneurysm). In contrast, a large maximum is obtained in the right side. Fig. 5.9(d) shows all maxima. A large part of the vessels is already removed, because they connect to the main large vessels with higher contrast.




Figure 5.9: Detection of microaneurysms (I). a) Original image, b) Result of pre-processing, c) Division by 2 and reconstruction, d) Maxima.

This method fails in an extreme case, for example, when a microaneurysm is too close from a large dark structure. This situation is illustrated in Fig. 5.10. In the center of the image, a microaneurysm is marked by the specialists. It is attached to a vessel and has low contrast. The reconstruction razes it. As a result, the microaneurysm is merged with the vessel.

A preliminary selection. A sequence of selections is applied to the obtained maxima to remove false positives. First, an area opening with criteria d_{MA}^2 is used to remove large objects (Fig. 5.11(a)). The mean contrast of each connected component is calculated on I_{pre} (the same method as in Section 4.5.1 is used). All connected components with a mean contrast less than 5 are removed (Fig. 5.11(b)). Elongated structures are removed by geodesic length (see Lantuéjoul and Maisonneuve [59], Morard et al. [70]). The criterion is set to $1.5 \times d_{MA}$ (Fig. 5.11(c)). Objects smaller than 6 pixels are regarded as dot noise and are removed by an area Opening. Meanwhile, candidates within the region of the




Figure 5.10: A microaneurysm over a vessel. a) Original image, b) Result of reconstruction

Criteria	True positives / False negatives	False positives
Maxima	1294 / 11	2,112,787
Area	1292 / 13	2,110,164
Contrast	1283 / 22	244,896
Geodesic length	1265 / 40	229,606
Dot noise and Optic Disc	1257 / 48	144,769

Table 5.1: Candidates selection. Number of TP, FN and FP after each selection step.

Optic Disc are removed (Fig. 5.11(d)). The obtained binary image is denoted $I_{candidate}$, which contains the binary mask of all candidates.

The selected criteria are all simply based on contrast and morphology. For each step, the number of true positives, false negatives and false positives are given in Table 5.1. As in the previous section, the experiment is done with the e-optha MA database, i.e. 148 images containing 1305 microaneurysms. Comparing with the annotations by the expert, the result of each step is analyzed. Based on the image maxima, 1294 microaneurysms are detected, together with two million false positives. The steep drop in false positives comes after the selection by contrast. More than 90% false positives are removed, while only 22 microaneurysms are lost. The final result contains 1257 true positives, which corresponds to 96.3% sensitivity. However, a large number of false positives (more than 140000) remains. On one hand, compared to the raw maxima, a huge amount of false positives are removed. On the other hand, the result is far from satisfactory. A finer analysis is needed to obtain a useful result. The result of the current detection will be considered as the set of microaneurysms candidates. In the following sections we will classify them using machine learning methods with appropriate features.




Figure 5.11: Detection of microaneurysms (II). a) Area opening, b) Thresholding, c) Thinning by geodesic length, d) Final candidates by removing small objects and candidates inside Optic Disc.

5.3.4 Features extraction

In order to remove the false detections of microaneurysms, more features should be extracted and analyzed. In this section, all selected features will be introduced, including several new proposed contextual features. A random forest is used to perform the classification. Then a multi-threshold implementation boosts up the classification result. The features can be summarized in three categories: intensity, geometric and contextual features. The same code rule of the descriptors of exudates is used here. Some descriptors are computed in the same way as for the exudates (Section 4.5.1).

Intensity features

The maximum and mean values of each candidate in the contrast image (I_{pre} , see Fig. 5.9(b)) provide basic intensity information. They are coded as *maxContrast* and *meanContrast* and computed by the same method used in the exudate descriptors computation. The median value of each candidate is

also computed, and the ratio between this value and $maxContrast$ is taken as a descriptor, coded as R_{median} . From the microaneurysm model introduced in Section 5.2, we deduce that this value should be around 0.5. If it is close to 1, it could be a flat object with sharp edge. On the contrary, if it is close to zero, it could be a low flat structure with a high impulse.

Geometric features

The descriptors $area$, $geoLength$ and $Circularity$ are calculated by the same method as for the exudates. The $area$ is normalized by d_{MA}^2 and $geoLength$ is normalized by d_{MA} . The descriptor $strain$ is given by the following expression:

$$strain = \frac{\min(H,W)}{\max(H,W)} \quad (5.1)$$

where H and W are the height and the width of each candidate.

Contextual features

Let us begin with two examples to illustrate our purpose. Fig. 5.12 shows a vessel fragment which has very low contrast (top right) and becomes two separated structures in the pre-processed image (middle right). Then they are detected as microaneurysm candidates (bottom right). If they are analyzed locally by contrast or shape, it is hard to distinguish them from microaneurysms. Another example has been previously shown in Fig. 5.7. In the green channel (left), the central spot is not perceived as a microaneurysm because of the high intensity in the surrounding area and the average background gray level throughout the image. But a simple local intensity analysis would lead to a wrong detection. In the work of Fleming et al. [30], the author deduced that the recognition of microaneurysms is dependent on the contrast in the surrounding retina. This is considered as a contextual feature.

We focus on the situation where the false detections are caused by the vessels. In the state-of-the-art, a vessel mask is obtained beforehand. Then a “hard” removal of the microaneurysm candidates is applied within the mask. However, the automatic segmentation of vessels is a hard task. Some microaneurysms and hemorrhages close to the vessels are easily segmented as vessels. We propose a new “soft” method to distinguish these two cases using Gabor filters. Fig. 5.13 shows two examples. Disconnected vessel fragments appear in the center of Fig. 5.13(a). A microaneurysm near vessels appears in Fig. 5.13(e). They are all segmented as vessels in the binary vessel mask (Fig. 5.13(c) and Fig. 5.13(g) can be practically regarded as a binary vessel mask). These two situations are schematized in Fig. 5.14. The disconnected vessel fragment and the true microaneurysm are all detected as microaneurysm candidates and segmented as vessels too (Fig. 5.14(c)). The microaneurysm candidate region is denoted P_{cand} (red). The binary vessel masks to be analyzed is denoted P_V (blue). The vessels within the region-of-interest is denoted $P_{V_{ROI}}$ (yellow contoured). The region-of-interest is selected by the vessel orientation map. First, all candidate pixels are added to $P_{V_{ROI}}$. Then, the orientation of the candidate (α_{cand}) is defined as follows:

$$\alpha_{cand} = \underset{\alpha}{\operatorname{argmax}} N_{\alpha} \quad (5.2)$$

where N_{α} is the number of pixels with angle α within the candidate (α is given by the orientation map obtained during the vessels segmentation, Appendix B). Then the connected neighbor pixels with an angle of $\alpha_{cand} \pm 30^\circ$ are added to $P_{V_{ROI}}$. This process is illustrated in Fig. 5.15. This region-of-interest




Figure 5.12: Disconnected vessels become false positives. Top right: green channel; Middle right: pre-processed image; Bottom right: microaneurysm candidates.

is very useful to find whether the segmented vessel is really a vessel fragment or is just a microaneurysm connected by the operator of vessel segmentation.

Four descriptors are proposed based on these observations. A window of width $6 \times d_{vessel}$ centered at the central point of the candidate spatially limits all the following calculations (green windows in Fig. 5.14). The mean contrast (on I_{pre}) within the region $P_V - P_{cand}$ and mean contrast within the region $P_{V_{ROI}} - P_{cand}$ gives information about whether the candidate is in a vascular zone, denoted $meanContrast_V$ and $meanContrast_{V_{ROI}}$. More importantly, if there is a real microaneurysm, the $maxContrast$ (of the candidate) will be much higher than $meanContrast_{V_{ROI}}$. In contrast, if there is a vessel segment, this difference is very weak. A comparison between Fig. 5.13(d) and Fig. 5.13(h) illustrates this observation. The ratio between $maxContrast$ and $meanContrast_V$ and the ratio between $maxContrast$ and $meanContrast_{V_{ROI}}$ are taken as two descriptors and coded as R_V and $R_{V_{ROI}}$, just for the sake of simplicity.

All the descriptors are summarized in Tab. 5.2. We have tested many other descriptors, for example, intensity features on other color channels or even more complex operators, tortuosity and textual




Figure 5.13: Disconnected vessel fragments (first row) and microaneurysm near a vessel (second row). a,e) Green channel, b,f) pre-processed image, c,g) Orientation map on segmented vessels, d,h) Pre-processed image intersected with segmented vessels.

Intensity:	$maxContrast, meanContrast, R_{median}$
Geometric:	$area, strain, geoLength, circularity$
Contextual:	$meanContrast_V, meanContrast_{V_{ROI}}, R_V, R_{V_{ROI}}$

Table 5.2: Descriptor list

information in a neighbor region. But experiments show that they are either much less competitive or give redundant information.

5.3.5 Multi-threshold boost up

We have discussed in detail the multi-threshold decomposition used in features extraction in Chapter 2. The implementation by Max-tree decomposition provides a possibility to efficiently analyze a contrast image layer by layer. Using this strategy, several descriptors can be applied to different layers. They are R_{median} , $area$, $strain$, $geoLength$ and $circularity$. For each connected component, the features are computed for the maxima, as well as the following layers, until the intensity level reaches 20. The statistics show that the contrast of microaneurysm is mainly inside the interval between 5 to 25. For the candidates whose contrast is less than 20, the rest of the features are automatically filled by an infinite value. The output vectors are aligned using the introduced top-bottom method.




Figure 5.14: Schematic illustration of vessel fragments and true microaneurysm. a) Disconnected vessel fragments, b) A true microaneurysm near a vessel, c) Contextual analysis.




Figure 5.15: Selection of region-of-interest by vessel orientation map. The red structure is the candidate. The yellow structures are selected regions of interest.

5.3.6 Learning, features importance and classification

With multi-threshold boost up, 106 features are extracted for each microaneurysm candidate. A machine learning method is used to perform a classification using this information. Learning is based on the candidates. Two classes, microaneurysms or other structures, are assigned to each candidate from the expert's annotation. The final number of candidates on 148 images is more than one hundred thousand (details in Tab.5.1). If we use the entire e-ophtha MA database, there will be 233 healthy images added. The number of candidates will increase. To balance the two classes for the learning, we can use up-sampling: duplicate the true microaneurysm set, until the number equals to one third of the number of total candidates. This is so called up-sampling. Another strategy is down-sampling the false positive set to a number equals to the number of the true microaneurysm set. But a lot of candidates are lost during the sampling, which will introduce a bias. Here we use the up-sampling method. The model is trained by Random Forest (Breiman [13]). The number of trees is set to 200 by considering the performance and the efficiency.

The features importance is computed during the generation of the trees in the random forest. Two tests are performed here: with and without multi-threshold boost up. All candidates from 148 images with up-sampling are used to perform this test. The results are presented in Fig. 5.16(a) and Fig. 5.16(b). Because there are more than one hundred features after multi-threshold expansion, only the top-20 descriptors are presented. The number in the end of certain features represents the layer

number in its multi-threshold version. Three feature categories are highlighted with different colors. We can conclude that:

- Contextual features, which are related to the vessels, play an important role in the classification. They rank 3rd and 4th before multi-threshold boost up, and are present in the top-20. This is probably because the main source of false positives comes from vessel fragments. The proposed contextual features differentiate their characteristics very well.
- *geoLength* and *area* are not very competitive, especially with multi-threshold boost up. In contrast, they are very useful for candidates extraction. From another point of view, all the candidates are selected with these two criteria, so it is normal that they don't provide further information.
- Layers 7 to 10, which correspond to the middle layers of each candidate, are the most important. These information can be used to reduce the features in future work.
- R_{median} is an interesting feature. It ranks last in the first test, but with multi-threshold boost up, it occupies 5 places in the Top-20. Note that R_{median} in the first test are computed on the candidates' binary masks, which are the bottom layers in multi-threshold decomposition. The change of the importance proves that for the same feature, computing on different supports does provide more information.

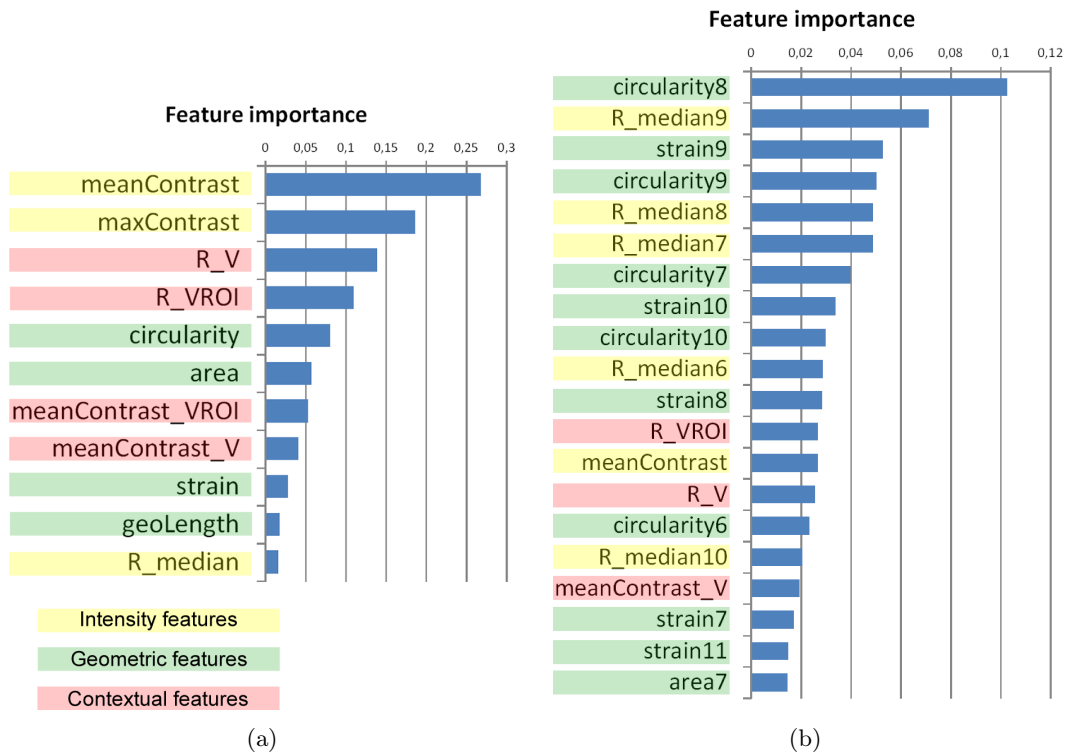


Figure 5.16: Features importance. a) Result without multi-threshold boost up, b) Result with multi-threshold and Top-20 features

The Random Forest gives a probability of being a microaneurysm to each candidate (Fig.5.17(b)). A large number of candidates are removed by assigning a very low probability. Details are shown in Fig. 5.18. Classification results and expert’s annotations are superposed on the original color image. The probabilities are mapped to gray levels between 0 and 255. Expert’s annotations are drawn by green circles. Four cases are zoomed in. From the top right, the first is a well detected microaneurysm. The classifier assigned a very high probability to this candidate. The second is the same case as in Fig. 5.10. A microaneurysm is attached to a vessel and it is not detected as a candidate. The third is a microaneurysm detected as a candidate. But the classifier assigned a very low probability to it, maybe because of its elongated shape and being near several other dark structures. The last case is a false positive. The classifier assigned a medium probability to it. It is at the end of a vessel and disconnected, but has a relative high contrast.




Figure 5.17: Classification result. a) Microaneurysm candidates, b) Result of Random Forest classification (probabilities are mapped to spectrum colormap), c) Ground-truth .

5.4 Method evaluation

By thresholding the probability of the candidates and comparing to the expert’s annotation, we can evaluate the global performance of the algorithm, for example, by drawing a ROC curve. Compared to exudates and hemorrhages, microaneurysms have a restricted size. It is reasonable to define a microaneurysm as a connected component. In other words, microaneurysms are countable. We first dilate the ground-truth set by 5 to reduce the potential error caused by manual annotation. Then a reconstruction from the dilated annotations under the candidates gives the True Positives (TPs). The remaining candidates correspond to False Positives (FPs). A reverse reconstruction from the candidate under the dilated markers followed by a subtraction from the markers gives the False Negatives (FNs) set. For individual candidates, it is not practical to define the True Negatives (TNs) set. The evaluation is performed at two levels: lesion level and image level.




Figure 5.18: Classification results superposed with expert's annotation (green contour). The gray level of each mark represents its probability as a microaneurysm. From top to bottom: 1) TP, detected with high probability. 2) FN, not detected. 3) TP, detected with very low probability. 4) FP, detected with medium probability.

5.4.1 Lesion level evaluation.

A method from Niemeijer et al. [73] is adopted to generate a FROC curve (Bunch et al. [14]) which plots the sensitivity against the average number of FPs per image. A final score is obtained by calculating the average sensitivities at selected points ($\frac{1}{8}, \frac{1}{4}, \frac{1}{2}, 1, 2, 4$ and 8). The test is performed on the 148 images containing microaneurysms from the e-optha MA dataset. A leave-one-image-out strategy is used. Three curves are plotted in Fig. 5.19 corresponding to three classifications with different sets of features. The blue curve is the result of using only intensity and geometry features. The green curve takes the proposed contextual features into account, too. The result shows a huge performance improvement. The red curve is the result of taking all features with multi-threshold boost up. The best score raised to 0.44. The sensitivities of the selected points are shown in Tab. 5.3.

According to the Retinopathy Online Challenge (details refer to Niemeijer et al. [73]), a database with 50 training images and 50 test images are available for a competition of microaneurysm detection algorithms. Those images have three types of resolution, from 768×576 to 1389×1383 pixels. The best score is 0.434 from the team drscreen (Antal and Hajdu [7]). Unfortunately, this database is no longer supported. We achieve a score of 0.44 with the same evaluation method, but of course this comparison has to be considered with precaution, as the database are different. However, we should note that we have 148 images containing microaneurysms and 233 healthy images in the e-optha MA database. Moreover, there are four types of resolutions, from 1440×960 to 2544×1696 in pixels. Database used in this work is larger, contains higher resolution images and is also more heterogeneous.




Figure 5.19: FROC curves for different sets of features.

False positives per image	$\frac{1}{8}$	$\frac{1}{4}$	$\frac{1}{2}$	1	2	4	8
Sensitivity	0.17	0.24	0.32	0.44	0.54	0.63	0.74

Table 5.3: Sampling of FROC curve.

5.4.2 Image level evaluation.

The previous evaluation gives an accurate but abstract performance of the algorithm. In clinical conditions for screening applications, it is enough to evaluate the results at the image level, which means presence or absence of lesions. Now we use the entire e-optha MA database (148 images containing microaneurysm and 233 healthy images).




Figure 5.20: ROC curve calculated on the whole e-optha MA database, using TOP-20 features, based on the presence or absence of MA in each image.

The maximum probability of all individual candidates inside one image is taken as the probability of the presence of microaneurysms for the image. The ROC analysis is performed. As the number of images is largely increased, in order to get a balance between the performance and the efficiency, we perform a feature selection at first. Two tests are run using the entire database to do a leave-one-out cross validation. Note that a strategy of sampling the candidates is used during the training phase to balance two classes (for microaneurysm candidates, false positives are much more than true positives). One test is using all 106 features after multi-threshold boost up, and another is using the top-20 features. The difference of the AUC of the ROC analysis is only 0.02. Thus, a final test using only the top-20 features, but a strategy of multiplying the candidates is done. The ROC curve is plotted

(Fig. 5.20). The corresponding Area Under ROC is 0.92. Note that the best AUC value of the tests on exudates is 0.95 on e-ophtha database.

5.5 Conclusions and perspectives

In this chapter, a statistical analysis has been done over the annotated microaneurysms from the entire e-ophtha database. It provides important characteristics, such as contrast and size distributions. Then, a new method for microaneurysms detection has been introduced. It can largely remove the FPs from the raw microaneurysms candidates, while keeping a high sensitivity. New features have been introduced for the classification. The new contextual features, using vessel information, is proved to be very useful. The application of the multi-threshold decomposition on certain features has improved the classification performance.

Though the ROC analysis has achieved a good result, the lesion level evaluation is still not satisfying. The classification result contains still too much FPs. In order to remove them, the local features extracted from the candidates themselves are very limited. The candidates should be placed in a larger scope, thus to study their relations to the neighboring structures and the environment. More contextual features should be studied. As we have used vessel related features, other similar contextual features concerning reflections and background noise can be added to the classification.

6.1 Introduction

Among the numerous studies dealing with eye fundus image analysis in the literature, only a few have treated the detection and segmentation of hemorrhages. This is mainly because researchers concentrate on microaneurysms and exudates, which most of the time appear earlier than hemorrhages during the progression of diabetic retinopathy. This is also due to the fact that hemorrhages are extremely difficult to detect, as we will see in the following section.

The detection of hemorrhages is important and useful. As previously said, in most cases they are not the first signs of diabetic retinopathy, but in some cases they can be the sole lesions present in the image. An accurate detection result provides additional useful information for the diagnosis system. Considering the great challenges and limited time dedicated for this study, the target is to develop an accurate detection algorithm. In other words, all hemorrhages might not be detected, but there should be a low number of false positives. Together with the detected microaneurysms and exudates, their detection is expected to improve the final decision of examination classification.

This chapter begins with the analysis on the main difficulties of the task at hand to give a clear idea of the challenges in this work. After a brief summary of the literature, a new hemorrhage candidate detection method is proposed. Then, several features are extracted, including a vessel related feature, which is new. The classification is obtained by Random Forest. The proposed method is evaluated on the e-ophtha HE database. Both lesion level and image level evaluations are performed. Errors are studied afterwards. In the end, conclusions and perspectives are discussed.

6.2 Main difficulties

Hemorrhages are dark lesions, as microaneurysms. They can be as small as microaneurysm, but they are usually larger, and can even cover a large proportion of the FOV. Their geometrical characteristics are similar to exudates, which can have any size and shape. They can appear anywhere on the eye fundus images, even across vessels. They can also cover the macula. Several examples of common hemorrhages are shown in Fig. 6.1. Dot hemorrhages (Fig. 6.1(a)) have sizes close to microaneurysms.

Fig. 6.1(b) shows a medium size but low contrasted hemorrhage. In Fig. 6.1(c), a large and low contrasted hemorrhage touches the vessel. Fig. 6.1(d) shows a complex case. A huge hemorrhage covers a vessel. There are also lesions of other types appearing beside hemorrhage. Flame hemorrhages (Fig. 6.1(e)) are a common type of hemorrhage, which are very difficult to detect, because they are similar to vessels. Fig. 6.1(f) shows a large area of hemorrhages.




Figure 6.1: Examples of hemorrhages. a) Small dot hemorrhage, b) Medium size hemorrhage, c) and d) Large hemorrhages connected to vessel (from HEI-MED database), e) Flame hemorrhage, f) Huge hemorrhage (from STARE database).

The difficulty of hemorrhages detection comes from their extreme variability in shape and contrast, which makes them difficult to detect, and easily causes false positives. Low contrasted hemorrhages (Fig. 6.1(b)) are always hard to detect. Extremely large hemorrhages, like the one in Fig. 6.1(f) need special attention too, because most candidate detection methods are size limited. False detections mostly come from isolated pieces of vessels (Fig. 6.2(a)). Sometimes, even inside continuous vessels, an area with slight dropping of intensity can be considered as a hemorrhage candidate. Another kind of false detections are caused by surrounding bright structures (Fig. 6.2(b)). A clearly visible macula (Fig. 6.2(c)), of course, can be detected as a hemorrhage, if no special care is taken.

There are two specially difficult cases: hemorrhages attached to vessels (Fig. 6.1(c)) and flame hemorrhages (Fig. 6.1(e)). Individually, they are not hard to be detected, because they are well contrasted and have clear contours. However, by trying to detect them, a lot of vessels will also be detected as hemorrhage candidates. If the vascular network is perfectly segmented, it will be much




Figure 6.2: Main sources of mis-detected hemorrhages. a) Isolated vessels, b) Valley inside bright structures, c) macula

easier to find hemorrhages. However, automatic and accurate segmentation of the vessels is very hard. Thus, it is necessary to find a trade-off between them. Moreover, we cannot simply remove segmented vessels to reduce the number of hemorrhage candidates.

A similar situation happens with hemorrhages in the macula region (Fig. 6.3). They are easy to detect. However, most algorithms have a separate step which removes the macula. Thus, these hemorrhages are probably regarded as macula and removed. Note that lesions appearing in the macula region are very dangerous, and are an important indication for diabetic retinopathy grading.




Figure 6.3: Hemorrhages in the macula region.

6.3 Hemorrhages detection method

6.3.1 State of art

In the literature, some related studies treat hemorrhages together with other red lesions. Kande et al. [53] used matched filters by rotating Gaussian shape kernels to enhance dark structures including red lesions. Then, a relative entropy thresholding is used to get candidates. Vessels are removed by a

Top-Hat using rotating linear structuring element. 12 features are extracted from the co-occurrence matrix and Support Vector Machine (SVM) are used to perform the classification. Pradhan et al. [77] used a similar processing scheme to detect all red lesions. An automatic seed generation method, which is based on local thresholding and region growing, is used to detect red lesion candidates. Then classification is done by a hybrid classifier using k-Nearest Neighbors (kNN) and Gaussian Mixture Models (GMM). A new feature, elliptic variance, is introduced and proved to be very useful to remove false positives. Similar methods can be found in García et al. [37] and Zhang and Chutatape [115]. Bae et al. [10] proposed a hemorrhages dedicated method. Template matching is used to detect hemorrhage candidates. The template is defined by an inner circle and an outer square to differentiate foreground and background structures. After thresholding, region growing is used to get the contour of the candidates. Then, six features (area, compactness etc.) are used to remove false positives. A Foveal filter is used.

Morphological methods have been often used in this context. Tang et al. [102] and Joshi et al. [52] proposed a splat based segmentation and feature classification method. A splat is a collection of pixels with similar color and spatial location. Splats are generated by a watershed over a multi-scale gradient image. Splat features are extracted from pixel-based responses (color, response of filters etc.) and splat wise (shape, orientation etc.). A kNN classifier is trained to perform the classification. Zhang and Fan [116] treated all types of lesions together, even exudates. H-maxima and H-minima are used to select high contrasted maxima and minima and keep them as seeds. Then a multiscale morphological processing based edge model generates a scale map. Entropy thresholding is used to get precise contours. The largest connected components in the dark lesions are considered as vessels and removed. Circularity is used to remove linear structures.

Some works emphasize the importance of vessel removal. Hatanaka et al. [45] proposed to segment hemorrhage candidates and vessels together by finite difference calculations with two mask scales. All the segmented structures connected to the optic disc are supposed to be vessels and removed. The remaining vessels are removed by a thinning technique. Then the length-to-width ratio is used do a further selection of the candidates. A number of features are extracted to perform a classification. Fleming et al. [29] used the supremum of Top-hats by multi-scale rotating linear structuring elements to get hemorrhages candidates. Then a region growing method is applied to complete the candidates. Local vessels are removed by a tangential shift operator and a SVM classifier. Features like shape, intensity, vessel likelihood etc., are extracted and another SVM classifier is trained to perform the final classification.

From these proposed methods and considering the properties of hemorrhages, we find several interesting points but also some limitations. On the one hand, methods based on constant size criteria, like template matching, are inefficient, given the multi-scale nature of hemorrhages. On the other hand, methods based on the morphological reconstruction tend to produce better results. Our approach is inspired by these methods. Concerning vessel removal, we think that a hard decision, for example, removing the largest connected component in the candidates and removing anything touching the optic disc, is not appropriate. This is because hemorrhages can be very large and cross vessels. We will see such examples in the following. Thus, a soft vessels removal method is proposed here.

6.3.2 Hemorrhages candidates detection

As discussed in the main difficulties section, it is barely possible to use size and shape information to detect hemorrhages. We propose a very general method to extract hemorrhage candidates. The method is mainly based on the fact that hemorrhages are darker than the background. The key point

is using a morphological reconstruction. Reconstruction does not need any size and shape information, but an appropriate mask. Then the difference between the reconstructed image and original image gives the candidates.

To define a mask, one possibility is using a morphological fill holes operator. In our case, it produces a large mask covering the entire field of view. Ideally, hemorrhages will be flooded and appear in the residual image. However, when the hemorrhages are inside a larger isolated “valley”, the entire region will be flooded when filling holes. A huge candidate containing background parts will be obtained. Another possibility is using a closing by reconstruction. But, since the size information is not available, it is difficult to define the size of the dilation. Another common method is to create the mask by adding a constant value to the original image. However, the contrast level differs between individuals. This method also has the same problem as the fill holes approach. We propose a hybrid method to create the mask, as shown in Fig. 6.4.




Figure 6.4: Flow chart of the hemorrhages candidates detection method.

The first mask is created by a background removal and thresholding. Usually, all hemorrhages are contained in this mask, except for two kinds of hemorrhages: huge and low contrasted hemorrhages. Because the background estimation is size related, only part of extremely large hemorrhages will be detected in the mask (see for example Fig. 6.1(f)). Low contrasted hemorrhages will also be missed. Thus, a second mask is computed by adding a constant value to the image, and removing certain vessels in the mask to avoid the “valley” effects. Then, a selection based on geodesic length and circularity is used to remove large vessels. In the end, candidates situated in the OD and Macula region are removed. These steps are described hereafter.

Background removal

We have proposed a background removal method based on alternating sequential filters in Chapter 2. We have successfully applied it to vessel segmentation and microaneurysms detection. The maximum size used in this method is the width of the vessels. In the case of hemorrhages, if we want to distinguish them completely from the background, a very large size would be needed which would excessively slow down the computation. Moreover, the advantage of the previous method is that it keeps lots of details. In the present case our objective is to get an approximate mask of the hemorrhages. Thus, the background will be estimated by a large mean filter.

Above all, several preprocessing steps are applied. Thin bright structures are removed by means of an inpainting of size $d_{vessel}/3$ (as used for the exudates detection, but start with an opening,

Section 4.3). Border reflections will raise the intensity level in their neighbor region in the background estimation. A border reflection segmentation method has already been introduced for the optic disc detection (Appendix C). Here we remove the segmented reflections regions before computing the mean filter. The zero background is replaced by the global mean value (on set) of the original image. A mean filter with window size d_{OD} is applied on the processed image to estimate the background. Then, the negative residue gives the dark structures, denoted I_{BGfree} . Fig. 6.6 shows the background removal with and without border reflections.




Figure 6.5: Bright structures removal by morphological inpainting. a) Original green channel, b) Inpainted image.




Figure 6.6: Influence of border reflections on background removal. a) Original inpainted image with border reflections, b) Background removal (negative residue) on original inpainted image, c) Inpainted image with border reflections removal, d) Background removal on border reflection free image.

A thresholding at 2 on the residue image gives the dark structures (Fig. 6.7(b)). Some border regions appear in the image, but there is nothing in the original image. This is because, while fitting the zero background by the global on set mean value, this value is higher than certain border regions.

It creates the same effects as the border reflection. But we will see that this kind of markers will not give any hemorrhage candidate in the following reconstruction step. Then, an area opening removes objects smaller than $d_{MA}^2/3$ pixels. Finally, a dilation of size $d_{vessel}/2$ (hexagonal structuring element) gives the first mask (Fig. 6.7(c)).




Figure 6.7: Thresholding to get first mask. a) Background removal (I_{BGfree}), b) Thresholding at 2, c) Dilation after removing small structures.

Reconstruction by raising global intensity

As for the computation of h-minima, we add a constant to the original image and then reconstruct the result over the original image. This fills the local basins regardless of their sizes. The residues give hemorrhage candidates. This is useful for detecting extremely large hemorrhages and low contrasted ones. For example, the hemorrhages in Fig. 6.8 are well detected by reconstruction. The added constant is defined by the maximum intensity level of the background removed image I_{BGfree} in the previous step. But if this value is too high, the limit of the constant is set to 50.

However, sometimes non-uniform illumination causes large isolated dark regions in fundus images (see for example Fig. 6.8(c)). The residues are shown in Fig. 6.8(d), which contain dark structures and dark background regions. One solution is to reduce the mask's intensity where we want to maintain the original image. But we don't need to reduce the mask's intensity on the entire region, but several points inside this region. The vessels are the perfect reference. They usually cover the entire image. For example, in Fig. 6.8(d), vessels go through both dark regions. Moreover, by removing the vessels in the mask, they will also be removed from the candidates of hemorrhages afterwards.

Vessel analysis and final mask generation

While using vessel information, we should always keep in mind that there could be hemorrhages mis-segmented as vessels and that there are also hemorrhages attached to the vessels. Thus, we should not do a hard elimination such as removing everything touching the vessels or the vessels themselves from the candidates. Thus, to modify the reconstruction mask, we won't reduce the intensity level on the entire vessel mask region, but select some pixels from the vessel mask, which are considered as the




Figure 6.8: Reconstruction by raising global intensity. a) and c) Original image, b) and d) Residues between the reconstructed and original images. Arrows in c) and d) indicate unwanted detections.

most likely to be vessels. The analysis is based on the circular variance. The result of this analysis will be not only used here, but also for the features extraction.

The objective of this analysis is to find the suspicious parts in the vessel mask. A binary vessel mask and its corresponding orientation map are needed. Any vessel segmentation methods could be used here. The orientation map gives the local orientation of the structure where the pixel is situated. Vessels segmentation and orientation map computation are detailed in B. The binary vessel mask and the orientation map are illustrated in Fig. 6.9(a) and Fig. 6.9(b). The orientation map is represented by different colors, in fact, each pixel p with a non-zero value in the vessel mask has an angle value $\theta_p = \frac{k\pi}{12}$, where $k \in \{0, 1, \dots, 11\}$.

If we look into the details of real vessel regions (Fig. 6.10(a) and Fig. 6.10(d)), the orientation map appears homogeneous. In contrast, those for the hemorrhages (Fig. 6.10(b) and Fig. 6.10(c)) are quite heterogeneous. Besides, vessels crossing or bifurcation areas show the same property as the hemorrhages. In circular statistics, circular variance can describe this difference (MacArthur and Thornton [62], Fisher [27]). A common definition of circular variance is given by:




Figure 6.9: Vessel analysis. a) Vessel mask of Fig. 6.5(a), b) Orientation map presented in color, c) Circular variance. d) Selected vessels by thresholding the circular variance at 0.25.

$$CircVar(p) = 1 - \frac{1}{N} \sqrt{\left(\sum_{i=1}^N (\cos \theta_{p_i})^2 \right) + \left(\sum_{i=1}^N (\sin \theta_{p_i})^2 \right)} \quad (6.1)$$

For each pixel p , the circular variance is computed within a moving window. N is the total number of the pixels used in the computation inside the window. This variance value lies in $[0,1]$. If all points have similar direction, the circular variance will be very small.

The window size is set to $2d_{vessel}$ in our case. Only non-zero points in the vessel mask participate in this computation. The results of circular variance on the previous examples (see third row in Fig. 6.10) show higher values in the hemorrhages regions. Keeping the pixels whose value is lower than 0.25 filters most big hemorrhages out from the vessel mask. The selected vessels are shown in Fig. 6.9(d), where hemorrhages and vessel crossings are removed.

To create the second mask, after adding to the original image a constant, the pixels which still remain after vessel analysis are set to their original intensity level. The first mask is set to the maximum intensity value (255). Then, the two masks are combined together by taking the supremum (Fig. 6.11(a)). The residue of the reconstruction is shown in Fig. 6.11(b). This image is called the




Figure 6.10: Details of vessel orientation map and circular variance. Top row: original images; middle row: orientation maps; bottom row: circular variance.

contrast image and denoted $I_{contrast}$. By looking at the original image (Fig. 6.5(a)) and the result of background removal (Fig. 6.7(a)), some vessels are removed and hemorrhages with details are kept. The results of another two examples (Fig. 6.8) are shown in Fig. 6.12. Note that the “valley” effects of reconstruction illustrated in Fig. 6.8(d) are corrected by removing vessels from the mask.

Length and circularity selection

A very limited selection is applied here to remove noise and other structures, because features extraction and machine learning will be used afterwards. Two kinds of objects are considered. Small objects which are treated in the scope of microaneurysms will be removed by geodesic length thinning. Extremely elongated structures are considered as vessels and removed by a circularity thinning. Since hemorrhages could touch the vessels, criteria thinnings will be applied to the gray level image. Neither the geodesic length or circularity are increasing criterion. After computing the criterion for each node in the max-tree, the “Max” decision is adopted (Salembier et al. [85]).

A closing with a single hexagonal structuring element is applied on $I_{contrast}$ to fill small holes which may affect the result of the following computation. Then, the geodesic length thinning is applied to remove structures shorter than $1.2d_{MA}$ pixels (Fig. 6.13(a)). The circularity thinning with criterion 0.14 is applied on the previous result and taking the residue (Fig. 6.13(b)). Note that 0.14 is a very




Figure 6.11: Combined mask and reconstruction. a) Final mask used for over reconstruction. b) Residue image I_{res} .




Figure 6.12: Examples of contrast image. a) Contrast image of Fig. 6.8(c), b) Contrast image of Fig. 6.8(a).

strict criterion. From the resulting image, we can notice that only extremely elongated structures are removed. In the end, after thresholding from 2 and removing the candidates inside OD and macula region, the **candidate image** is obtained and denoted I_{candi} . Several examples of extracted hemorrhage candidates are shown in Fig. 6.14. We can see that hemorrhages from small to large sizes are detected. Those with low contrast and attached to vessels are also detected. Large vessels are removed, though small pieces are still kept as candidates. In the following features extraction step, we propose special features to deal with this problem.

6.3.3 Features extraction and classification

We extract 7 representative features for each candidate, which can be summarized in three categories: contrast, geometric and vessel likelihood features. Our experience on exudate and microaneurysm fea-




Figure 6.13: Candidate selection. a) Geodesic length thinning, b) Circularity thinning.

tures extraction tells us that basic information, such as contrast and geometry, always plays important roles in the classification. Special features such as the distance to the nearest vessel for exudates are sometimes very useful too. From the detected hemorrhage candidates, we can see that the main source of noise comes from the vessels. We will present a new feature, which is called vessel likelihood and show its interest.

Since most features are easy to be computed, and have already been introduced in the exudate and microaneurysm chapters, we will give a brief summary here. There are three features based on the contrast:

- Mean value on the contrast image (*meanContrast*).
- The ratio between the mean value on the original green image within the candidate and its neighbor region (*R_meanGreen*). The neighbor region is defined by the same method as used for exudate detection, expanding the candidate's bounding box by d_{vessel} in each direction and removing the candidate.
- A similar computation as *R_meanGreen*, but we compute the difference between the candidate and its neighbor region rather than the ratio (*D_meanGreen*).

Three geometric features:

- Area (*area*).
- Geodesic length (*geoLength*).
- Circularity (*circularity*).

A vessel likelihood feature was also used by Fleming et al. [29]. But the definition is totally different. In their work, a support vector machine is trained to classify the vessels. The maximum probability of the vessel classification result inside the hemorrhage candidate is taken as vessel likelihood. In our case, firstly, we don't have a soft classification step for the vessel. The circular variance (Eq. 6.1) computed on the vessel mask is a good feature for distinguishing vessels from hemorrhages. Secondly, it is not appropriate to simply take the maximum probability of the vessel likelihood inside a hemorrhage




Figure 6.14: Examples of final hemorrhage candidates. a), b) and c) Original green image. d), e) and f) Hemorrhage candidates.

candidate, because there are many cases where the hemorrhages are attached to the vessel. If the candidate mask contains both a hemorrhage and a vessel, the maximum probability of vessel likelihood will be very high. For example, Fig. 6.15 shows a hemorrhage attached to a vessel. In the candidate mask, part of the vessel is segmented as hemorrhage. From the circular variance image (Fig. 6.15(c)), we can see that some parts have very low value. In other words, if we only take the maximum vessel probability, this hemorrhage will be likely classified as vessel.

We propose to compute the ratio of the weighted vessel area over the candidate area as the vessel likelihood. The main idea is to count the number of vessel pixels inside each candidate. If most pixels inside a hemorrhage candidate are in the vessels mask, this candidate is probably a part of vessels. A lower ratio value indicates a higher probability of being a real hemorrhage. However, sometimes, hemorrhages are segmented as vessels (Fig. 6.15(c)). In this case, the ratio value is very low, if we directly count the number of the pixels inside the vessels mask. In order to minimize the influence of inaccurately segmented vessels, the vessels mask is assigned weight values according to two criteria.

The first criterion is the circular variance. The same threshold at 0.25 is used. If the value of circular variance is lower than 0.25, the pixel satisfies the first criterion ($C1$). The second is the vessel length. Though we have exceptions for huge hemorrhages and hemorrhages attached to vessels, most




Figure 6.15: Vessel likelihood feature. a,e) Original green image, b,f) Detected hemorrhage candidate, c,g) Circular variation computed on vessels mask, d,h) Selected vessels by circular variance (red).

hemorrhages are detached from vessels after all. We say that if the geodesic length of a vessel is larger than $d_{FOV}/3$, all the pixels satisfy the second criterion ($C2$). Thus, if a pixel from the vessel mask satisfies both criteria, it is probably a real vessel. While computing the vessel area inside a candidate, the weight of the pixels is defined as follows:

- If a pixel satisfies $C1$ and $C2$, it is counted as 1 pixel.
- If a pixel satisfies $C1$, but not $C2$, it is counted as 0.7 pixel.
- If a pixel satisfies $C2$, but not $C1$, it is counted as 0.3 pixel.
- If a pixel satisfies neither $C1$ nor $C2$, it is counted as 0.1 pixel.

In the end, the ratio between the vessel area and the candidate's area (LH_vessel) is taken as the vessel likelihood feature. this value is between 0 and 1. As for the previous example, a real hemorrhage (Fig. 6.15(a)) has a vessel likelihood value of 0.3. However, the candidate in Fig. 6.15(f) has a value of 1. The difference is huge. This is the advantage of making a soft decision on whether a candidate is a hemorrhage or a piece of vessel. If not, they will be removed together.

6.3.4 Classification

A supervised classification is given by Random Forest (Breiman [13]), the same as for the classification of exudates and microaneurysms. The number of trees is set to 500. The features importance are shown in Fig. 6.16. The *area* ranks first, which is reasonable for hemorrhages. After excluding the vascular network and the macula, a large size candidate is probably a hemorrhage. The other features are almost equally important. The new proposed vessel likelihood feature (LH_vessel) ranks slightly higher than the *meanContrast*, which means it is useful for the identification of hemorrhages.




Figure 6.16: Features importance.

The result of the classification is illustrated in Fig. 6.17. Note that in this case the performance seems very satisfactory. The method is quantitatively evaluated in the following section.

6.4 Method evaluation

6.4.1 e-ophtha HE database

In the e-ophtha database, three types of hemorrhages are annotated by the expert: dot, blot and flame hemorrhages. Dot hemorrhages have small sizes. Flame hemorrhages have elongated shapes. The rest are blot hemorrhages which have various sizes and shapes. There are 99 images containing annotated hemorrhages. In order to construct the e-ophtha HE (hemorrhage) database, first of all, we remove small hemorrhages which have already been merged to the e-ophtha MA database. Now, there remains 90 images containing 299 hemorrhages. Then, the 233 healthy images in the e-ophtha MA database are merged with the hemorrhage images to complete the database. We finally obtain 323 images in the e-ophtha HE database.

6.4.2 Evaluation

Two kinds of tests are used to evaluate the proposed algorithm on the e-ophtha HE database. The lesion level FROC analysis and the image level ROC analysis.

Lesion level FROC analysis

In the case of exudates, one large exudate could be accompanied by a bunch of tiny exudates. This fact makes it inappropriate for an evaluation by only counting the number of lesions. However, for hemorrhages, the situation is different. Though their sizes vary a lot, we detect small hemorrhages as microaneurysms. Therefore, there won't be small regions scattered with a bunch of hemorrhages. Thus, we will simply count the number of lesions for the lesion level evaluation. Besides, we consider that a hemorrhage is correctly detected (TP), if it touches the ground-truth.

The FROC curve plots the sensitivity against the average number of false positive detections per image (Bunch et al. [14]). A leave-one-out (LOO) test is applied on the entire e-ophtha HE database.




Figure 6.17: Random Forest classification result. a) Original image, b) Manual annotation, c) All candidates, d) Probability map.

The FROC is shown in Fig. 6.18. The curve shows that there are 30% hemorrhages not detected even with the lowest threshold on the classification result. We did an analysis on the missed hemorrhages:

- As predicted, many flame hemorrhages (Fig. 6.19(a)) are missed, given their similarity with vessels.
- Hemorrhages appearing in the macula region (Fig. 6.19(c)). In this case, the hemorrhages are detected as candidates. But they are removed because they belong to the macula. We mentioned this problem in the beginning while analyzing the main difficulties. This problem is difficult to solve. Firstly, we should remove the macula, or it will always be classified as hemorrhage with high probability. Secondly, in contrast, hemorrhages in the macula region are very dangerous. To solve this problem, a deep study on the detection of the macula is necessary, especially on how to differentiate the hemorrhages from the macula.
- Hemorrhages with low contrast (Fig. 6.19(d)). This is the main source of FNs.
- Small hemorrhages. Some hemorrhages have a critical size between microaneurysm and hemorrhage. They may be detected as microaneurysm.




Figure 6.18: FROC analysis.

On the other hand, the curve shows that the number of FPs is low. This meets our original expectations, that is developing an accurate hemorrhages detection method, which may not detect all hemorrhages. However, we studied the main sources of FPs:

- The macula when its automatic detection fails.
- Low intensity regions inside large reflections (Fig. 6.20).
- Vessels. Although good features plus machine learning have removed a large amount of candidates from vessels, the confusion between them is not completely eliminated.

Image level ROC analysis

Image level evaluation is more interesting from a clinical perspective, especially when dealing with screening applications. Images are classified into one of two categories: presence of hemorrhages or absence of hemorrhages. The ROC curve on the e-optha HE database is shown in Fig. 6.21.

The result is very interesting. Since the detection of hemorrhages is a difficult task, and the algorithms are designed for an accurate but not overall detection, achieving 0.9 AUC is quite encouraging. Note that the best AUC values of the same image level tests on exudates and microaneurysms detection are respectively 0.95 and 0.92 on the e-optha databases.

6.5 Conclusions and perspectives

The problem of hemorrhages detection in eye fundus images has been analyzed in this chapter. A novel hemorrhage segmentation method has been introduced. This work is among the first to treat this problem.




Figure 6.19: Examples of false negatives. a) and b) Flame hemorrhages, c) Hemorrhages in the macula region. d) and e) hemorrhages with very low contrast. Red circles indicate the experts annotations which have been missed by the algorithm.

One contribution of this work is the special designed mask used to extract hemorrhage candidates. It is efficient regardless of size. Another important contribution is the use of vessel information, based on the circular variance, to distinguish between vessels and hemorrhages. Vessel information is used for mask generation and feature extraction. Unlike previous approaches, which remove hemorrhage candidates attached to vessels, we propose here a soft approach.

The lesion level and image level evaluations on the e-optha HE database show encouraging results. A 0.9 AUC value is achieved in the ROC analysis, with 90% sensitivity and 73% specificity. Behind these good results, 30% of hemorrhages are not detected. Analysis of the method drawbacks open the way to several possible improvements:

- A more accurate macula detection method. As mentioned several times, the removal of the macula is important but dangerous.
- In spite of the initial suppression of thin bright structures applied during the background removal phase, some candidates are introduced by spurious bright structures. The method proposed by Zhang and Fan [116], based on a local histogram analysis, could be applied here to improve our results.
- A further study on flame hemorrhages. Flame hemorrhages represent an important part of hemorrhages. Their detection should be linked to vessel segmentation. Flame hemorrhages are isolated, usually highly contrasted. A vessel tracking method may be useful. For example, when an isolated vessel is found, it could be compared to the nearest vessel based on contrast, intensity, width or edge sharpness. Another possibility is to consider them as a third class in the classification.




Figure 6.20: Examples of false positives. Low intensity region inside large reflections. a) Hemorrhage candidate with high probability, b) Original image.




Figure 6.21: ROC analysis and area-under-curve (AUC)

This method has not yet been integrated into the TeleOphta system. It will be interesting to evaluate its contribution to the screening system.

The diagnosis of diabetic retinopathy is performed not only based on retinal images, but also on contextual information, such as the patient age, weight or diabetes history. The TeleOphta system classifies the examinations acquired in a teleophthalmology network into one of two categories: “To be referred” or “Not to be referred”. The examinations contain color eye fundus images, as well as patient data.

7.1 Heterogeneous information fusion

The strategy used in TeleOphta takes all the factors in the examinations into account (illustrated in Fig. 1.2, Chapter 1). The heterogeneous information is fused to generate the final classification. This part of the work was developed by G. Quéllec, M. Lamard and G. Gazuguel (LaTim - Inserm U1101) and integrated by R. Danno, D. Elie and B. Laÿ (ADCIS). A brief introduction of the factors is presented as follows:

- Signature-based pathological score. It is obtained by a multiple-instance learning and weakly supervised classification method (Quéllec et al. [82]). First, a reference dataset is built up. Each image is categorized as relevant (containing diseases) or irrelevant by clinicians. Then, images are divided into multi-scale patches. The scales are adapted to different pathologies. Patches are described using textural and shape features. Zernike moments (Khotanzad and Hong [56]) and wavelet-based features (Quéllec et al. [79]) are used for shape and texture characterization. This feature vector is called a signature. A local relevance score is computed for each patch based on the signature, by comparing to the signatures of all other patches in the reference dataset. (Quéllec et al. [80]). Finally, a global relevance score is computed by combining all local relevance scores. This score is so called signature-based pathological score.
- Image quality. Mathematical morphology based image quality metrics are computed, for example, the average intensity of the morphological gradient and the average intensity in the residual image of alternating sequential filtering.

		TeleOphta / Expert	
		Exams to be referred (positive)	Normal exams (negative)
Ground truth	Exams to be referred (positive) 24%	19.2% / 19.4%	4.8% / 4.6%
	Normal exams (negative) 76%	24.3% / 14.1%	51.7% / 61.9%

Table 7.1: Confusion matrix of the results on e-ophta database at 80.9% sensitivity by TeleOphta and human expert.

- Pathological score per lesion detector. A joint Cumulated Distribution Function is computed on the previously extracted lesion size and probability. After a linear discriminant analysis, the result is mapped to a single pathological score (Quellec et al. [81]).
- Demographic information fields. For each patient, basic information is provided with the record, for example, age, sex and weight, etc.
- Diabetes-related information fields. The related physiological measurements such as blood pressure and cholesterol, etc. are also considered in the classification.

7.2 Results

The entire e-ophta database was divided into two parts. The first half of the database was used for training the fusion model. The test was performed on the other half. The output of the system on each examination is a score which can be interpreted as the probability of “To be referred”. By thresholding this value, a ROC curve was computed to evaluate the performance (red line in Fig. 7.1). A subset of 500 examinations were randomly chosen, on which double reading was performed by ophthalmologists. Supposing that the original annotation of the database is the ground truth, the sensitivity and specificity of the double reading could be computed (blue star in Fig. 7.1). For the same sensitivity as the expert (80.9%, blue star), the system obtained a specificity of 68%, which is lower than the 81.5% of the expert. However, it means that the system can remove about 70% normal or “Not to be referred” examinations in the OPHDIAT network. Considering that about 74% of the cases in the network are “Not to be referred”, this result can save half of the time that ophthalmologists spend on screening. This also means that 2.3 times more patients can benefit from the service. If we raise the sensitivity to 90%, a specificity of about 35% is obtained. This is also a very interesting result. It performs a more secure selection, while helping the experts to save up to 35% time.

The confusion matrix of both TeleOphta and human expert at 80.9% sensitivity is shown in Tab. 7.1. Because the points are chosen for the same sensitivity, the performances of positive cases detection are quite similar. However, one can confidently suppose that the errors made by the expert and by the TeleOphta system are not of the same nature. For example, human experts may have different opinions on an image which shows little pathologies, but the probability to ignore a severe or an urgent case is very low. However, for the system, it remains unknown.




Figure 7.1: ROC curves on e-optha database (red) and on recent database with 1066 images (green), and the performance of a specialist (blue star).

A second test was performed on a more recent database, in order to evaluate the stability of the classification quality with time. All 1066 examinations gathered by OPHDIAT in November 2012 were kept, after anonymization, in order to build this new database. It is possible to do an image level study for the wrongly classified cases. The classifier was trained on the e-optha database, and tested on the recent database. The corresponding ROC curve is shown by the green line in Fig. 7.1. With the same sensitivity as the human expert at 80.9%, we obtained about 50% specificity. It is still an interesting result, because half of the normal examinations are removed by the system, which means that about 1.8 times more patients can benefit from the diagnosis service. The confusion matrix at 80.9% sensitivity is shown in Tab. 7.2. It is reasonable that the performance on the recent database is poorer than on the e-optha database. The training of the TeleOphta classification system was performed by using the images of 2007 and 2008, while the test is performed on images of 2012. The OPHDIAT network evolves all the time. For example, new fundus camera models are used and the technicians in charge of taking the photos may change. Therefore, this means that we should update the training of the system regularly.

The 4.4% wrongly classified positive examinations (false negatives) correspond to 46 images. The grading provided by the experts indicate that among these three examinations were considered as urgent and one examination as very urgent (Fig. 7.2). In the images, signs of diabetic retinopathy

		TeleOphta	
		Exams to be referred (positive)	Normal exams (negative)
Ground truth	Exams to be referred (positive) 22%	17.6%	4.4%
	Normal exams (negative) 78%	39%	39%

Table 7.2: Confusion matrix of the results on the recent database (1066 examinations) at 80.9% sensitivity by TeleOphta.

such as microaneurysms and hemorrhages appear, but the reasons of urgency are different. These 4 urgent or very urgent cases were shown to two ophthalmologists. They confirmed the extreme urgency of the last case: one image shows what seems to be a papilledema, as evidenced by the blurring of the optic disc margins (Fig. 7.2(a)). They commented that this is an exceptional occurrence. This interesting case illustrates one of the main hurdles of computed assisted diagnosis system: the detection of exceptional cases. Fig. 7.2(b) and Fig. 7.2(c) are two cases of pregnant women. In this case, urgency is automatically increased. The fourth case (Fig. 7.2(d)) seems to correspond to a grading error: the two experts did not see any signs of urgency.

7.3 Conclusions

The final classification strategy and results of TeleOphta are presented in this chapter. A heterogeneous information fusion strategy is introduced, which combines the result of image processing and other contextual information. The output of the system assigns a score to each file, which makes it flexible to tune the system to satisfy different demands. Evaluations have been done on two large databases. The comparison between the system and human expert shows that, by adjusting the threshold on the final score, the system can achieve the same sensitivity as the specialist and a comparable specificity. With the same sensitivity as the expert, the result on the e-ophta database shows that 2.3 times more patients can benefit from the diagnosis service. Meanwhile the result on the database with 1066 recent examinations shows a slight descent for the specificity, but still 1.8 times more patients can benefit from the service. The performance of the TeleOphta system is already very interesting in this stage.

However, some exceptional urgent cases are hard to detect by a computer aided diagnosis system like the TeleOphta system. Several perspectives are presented to solve these problems:

- Give the exceptional cases to the machine learning module. However, establishing an appropriate list of exceptional examples is an extremely difficult task.
- Develop specific methods to detect certain cases, for example, analyze the optic disc. However, this is a hard work. It is difficult to exhaustively detect all possible exceptional cases.
- Develop a temporal analysis module of the patient records to automatically find the differences between two visits.




Figure 7.2: “urgent” and “very urgent” cases classified as “normal” by the TeleOphta system. The reason of urgency: a) Signs of papilledema to be confirmed, b) and c) Not urgent but pregnant women, d) Unknown but double reading downgrade it to not urgent case.

- Mechanisms to regularly update the training of the system have to be designed.
- Furthermore, it would be interesting to assess the importance of each piece of information in the final classification. Indeed, the TeleOphta system takes into account heterogeneous information, we do not know yet what kind of information is more important.

In the future, more wrong classified cases should be studied. The training should be done regularly, because fundus images change with the introduction of new fundus cameras. Temporal following is another important way to improve the results. By comparing the records of one patient, it is easier to find abnormalities.

Conclusions and perspectives

The main objective of this work is to develop automatic image processing methods for diabetic retinopathy lesions detection under clinical condition. This work is integrated in the TeleOphta project which aims at automatic detecting normal exams, so as to improve telemedicine networks for diabetic retinopathy screening. Classical methods usually fail when facing a heterogeneous clinical database. New methods are conceived in this work, during each step of fundus image processing, which not only solve the initial problems, but also adapt to the complex clinical conditions. A summary of the contributions and perspectives is given in this chapter.

8.1 Contributions

8.1.1 Contributions to fundus image processing

In order to achieve the final objective, i.e. accurate detection of diabetic retinopathy lesions, this work is divided into a number of sub-missions. A new gray-scale normalization method using residues of an alternating sequential filter and a new spatial calibration method which takes the width of field of view as invariant, are proposed. They make the following algorithms work on images of any resolution and with heterogeneous illuminations, while maintaining high performances.

This work concentrates on the detection of three lesion types: exudates, microaneurysms and hemorrhages. Some morphological approaches are for the first time applied in this domain, such as the use of the recently proposed operators ultimate opening and geodesic length thinning. These methods serve image filtering, candidates detection and features extraction, which is followed by machine learning to obtain more accurate results. Meanwhile, the solutions to specific problems during the detection of each lesion are important contributions, too. For example, the various types of reflections, commonly appearing in fundus images under clinical condition, are detected and removed for a better exudates detection result.

Contextual information provided by vessels has been studied in this work. Firstly, the density, width and orientation of the vessel mask indicate important location information of optic disc and macula. Secondly, they are used to segment reflections along the main vessels by a parabola fitting, which has been proved very useful for the exudates detection. Simultaneously, the distance from

exudate candidates to the nearest vessel has been shown to be an interesting descriptor. Moreover, given that vessel fragments are the main interferences of microaneurysm and hemorrhage detection, solutions by local contrast and directional analysis show good results on the removal of false detections, even using an imperfect vessel segmentation mask.

Though the detection and segmentation of retinal structures (vessels, Optic Disc and Macula) are not the final target of this work, they are still important intermediate steps. A fast vessel segmentation method based on residues analysis gives a proper vessel mask satisfying most requirements in this work. Meanwhile, another Gabor filter based method provides finer vessel segmentation and richer vessel information for a contextual analysis. A hybrid Optic Disc detection method and a preliminary Macula detection method provide necessary information for the major tasks.

Most proposed methods have been integrated into the TeleOphta classification system. Extensive tests show that this system already achieves an interesting performance, as compared to human specialists.

8.1.2 Contributions to general image and features analysis

Residues analysis by alternating sequential filters shows great interest used as preprocessing of fundus images. The proposed residual operator separates dark and bright parts of images according to texture and contrast information. Moreover, it corrects non-uniform illumination in images. In this work, it is applied to vessel segmentation and exudates and microaneurysms detection. It could be extended to the segmentations of any object with limited size in images with complex background.

Based on the advantages of residues analysis, a multi-threshold analysis is proposed to extract lesion features on gray level image with multi-supports for classification. It provides extra information of feature evolution according to different gray levels. It is particularly interesting when combined to the residues analysis, because the residues analysis set the start level of the contrast image at zero, which facilitate the alignment of multi-features. A Max-tree decomposition implementation makes it even more efficient.

8.1.3 Publications

The presented work has been the object of the following publications:

- Patent: Procédé de normalisation d'échelle d'images ophtalmologiques (Zhang et al. [119]).
- Medical Image Analysis: Exudate detection in color retinal images for mass screening of diabetic retinopathy (Zhang et al. [121]).
- IEEE International Symposium on Biomedical Imaging (ISBI): Spatial normalization of eye fundus images (Zhang et al. [120]).
- International Congress of Stereology (ICS): Application of the morphological ultimate opening to the detection of microaneurysms on eye fundus images from a clinical database (Zhang et al. [117]).
- The Association for Research in Vision and Ophthalmology (ARVO): Automatic detection Of exudates in color retinal images (Zhang et al. [118]).
- Ingénierie et Recherche Biomédicale (IRBM): TeleOphta: Machine learning and image processing methods for teleophthalmology (Decencière et al. [22]).

Moreover, two publications are under preparation:

- Hemorrhage detection.
- Building top-hat extensions using alternating sequential filters.

8.2 Perspectives

Profound studies of fundus image processing have been done in this work. Though good results have been achieved by the proposed methods, many aspects can still be improved. Since details have been discussed in each chapter, here we list some important perspectives.

The methods of retinal structures detection and segmentation could be improved. The results have impacts on the accuracy of lesion detection, e.g., removing the candidates located inside the segmented retinal structures increases directly the detection accuracy. On the other hand, using the contextual information provided by the retinal structures can help lesion detection. An accurate segmentation of these structures could help to detect other important structure related lesions such as venous abnormalities, or even other diseases like glaucoma.

As for the lesions detection methods, they usually fail on special cases, for example, those with the presence of large irregular reflections or laser treatment scars. These cases could be studied to develop special methods to increase the detection accuracy. For the detection of hemorrhages, though the proposed methods achieved a good score in global evaluation, there still remains some problems. The detection of flame hemorrhages needs to be improved with priority.

The multi-threshold analysis for features extraction is a good start. However, the raw features from different supports are just aligned together and passed to the classification step. However, they could be processed to reveal more implicit relations and differences during the features evolution, thus to improve the classification results.

The perspectives for TeleOphta project are linked to several aspects:

- All the developed methods, especially the lately proposed method of hemorrhages segmentation, need to be integrated in the final patient records classification system. The performance of the proposed methods in the final classification system need to be evaluated in detail. In other words, how do the results of detected lesions would affect the final decision should be studied. It would be useful for improving the segmentation methods and the classifier.
- Development of detection methods for other lesions. Lesions characterizing severe cases should be prioritized. For example, the detection of an abnormal optic disc could help detecting glaucoma cases, or papilledemas.
- Develop a temporal analysis method to measure the differences between two visits. This is an interesting and expected to be efficient solution to improve the overall performance of the TeleOphta system.
- From an application point of view, TeleOphta technology could be deployed outside traditional telemedicine networks. It could be adapted to images acquired with hand-held fundus cameras, or even with smart phones equipped with optic adapters. This evolution would bring the screening service to populations which are outside the perimeter of classical medical care.

Field of view segmentation

The mask of the FOV is used for spatial calibration and to delimit the region-of-interest of other processes. Two fundus images and their FOV binary mask are shown in Fig. A.1. In the literature, this work is omitted, because the background value is assumed to be zero and the FOV is much brighter. But in real conditions, the background is not zero. Moreover, in some badly illuminated images, there can be a very dark region in the border of the FOV, where the FOV and the background merge together. Fig. A.2 shows such an image. The red channel, which is supposed to be the brightest channel for fundus images, is in the same situation (Fig. A.2(b)). In this image, the pixels where the gray level value is equal to 0 or 1 are set to red and green. From the details (Fig. A.2(c)), we can see that pixels inside FOV have 0 and 1 intensity. Fig. A.2(d) shows the same image with a thresholding from 0 to 4 (included). Thus, a thresholding will not give a good segmentation of the FOV in this case.




Figure A.1: Examples of segmented FOV.

This section will propose a method to segment the FOV, which solves two problems: Automatically finding the proper threshold for different images to get the principal part; removing noise attached to the edge of the FOV.

Although the red channel is the brightest one among the three channels, in dark regions, they have the same low intensity levels. Thus, we take the supreme of red and green channels to start the process, denoted I_{sup} . The blue channel is excluded because it is too noisy. For a normal image, like the one shown in Fig. A.3(a), thresholding at 8 gives a good mask with low background noise. But

APPENDIX A. FIELD OF VIEW SEGMENTATION




Figure A.2: Example of bad quality image. a) Original image, b) Red channel (red and green color for intensity 0 and 1), c) Details, d) Thresholding from 0 to 4 (included).

in a difficult case like Fig. A.3(d), thresholding at 8 cannot get the entire FOV. But thresholding at 4 can give a reasonable result (Fig. A.3(f)). However, the same threshold on the previous image gives a result with too much noise (Fig. A.3(c)).




Figure A.3: Background noise level differs between images. a), d) Supremum of red and green channels, b), e) Thresholding at 8, c), f) Thresholding at 4.

In order to automatically define a threshold, firstly, an alternating sequential filter (ASF) starting with a closing is applied on I_{sup} to fill the holes inside the FOV and remove small variations. The size is set to 4 for all images, because the background "pepper and salt" noise has fixed small size. The resulting image is denoted I_{ASF} , its details are shown in Fig. A.4(b). Compared to I_{sup} , all the small holes in the dark region inside the FOV are filled.




Figure A.4: Result of ASF (red and green color for intensities 0 and 1). a) Details of I_{sup} , b) Details of I_{ASF} .

An analysis based on the area of the flat zones of I_{ASF} is used to find the proper threshold. Flat zones are displayed in colors in Fig. A.5(a). It is clear that the background is the largest flat zone after filtering by ASF. We define a minimum area criterion: $0.03 \times W \times H$, where W and H denote the width and the height of the image. Every flat zone satisfying this criterion is kept. The remaining zones are labeled by their intensity in the I_{ASF} . Then, the maximum intensity of the largest flat zone will be taken as the threshold. If no flat zone satisfies the area criterion, the intensity of the largest flat zone will be chosen as the threshold. In the illustration, the threshold is 3 (Fig. A.5(b)). From the details in the image, there is still some noise and the contour is not smooth. Sometime there are small holes inside the thresholded image. Thus, a morphological filling holes operation is applied to this binary image and the largest object is selected as the principal part of FOV, denoted I_{FOVP} .




Figure A.5: Threshold defined by the largest flat zones. a) Flat zones plotted in colors, b) Thresholded image.

The morphological Toggle-mapping (Meyer and Serra [68]) permits to compare the difference between an extensive function and an anti-extensive function and take the closest one. This operation is used to classify the pixels in the edge region, thus smooth the contour. Here, we use an erosion and a dilation with an hexagonal structuring element with size 5. First, the non-zero part of the binary mask is replaced by I_{ASF} . The pixels at the edge will be classified to background or to FOV by selecting the class with closer intensity. For example, Fig. A.6(a) is the edge region of I_{ASF} restricted by the mask. The intensity level between 1 and 4 are painted in red. The result of Toggle-mapping (Fig. A.6(b)) reduces the red region. Then, the result is thresholded from 1 and the biggest region is chosen as the new FOV. The whole process is repeated until the results are identical. In practice, it is usually under 5 iterations (a little more for big images, since we use a fixed size of structuring element). The image after 4 iterations is shown in Fig. A.6(c). The result of FOV segmentation of the example (Fig. A.2(a)) is shown in Fig. A.7. The dark zone is totally kept and has a smooth contour.




Figure A.6: Toggle-mapping is used to smooth the edge.




Figure A.7: Result of FOV segmentation.

An automatic FOV segmentation method is proposed in this section. It can deal with difficult cases which have low contrast and non-uniform illumination. This method also gives a smooth contour of the segmented mask. The proposed method is visually validated on a subset of e-ophtha database (about 1000 images). The FOV is correctly segmented. However, improvements by using shape information could be studied in the future. It is useful when a large dark zone is equal to zero, as it will be lost

by our method. Operations like Hough transformation could be used to reconstruct the border. The symmetry characteristic of the FOV can also be used, because the dark zone only appears on one side.

Vessel segmentation in color eye fundus images is interesting for several reasons. Vessels can help to identify retinal anatomical elements such as the optic disc and the macula. They can also provide important information to estimate some characteristics of the image, for example, Fleming et al. [31] used the visibility of the vessels in the macular region as an indicator of image quality. The accurate measurement of vessels can also help to diagnose some pathological conditions, such as venous beadings and intraretinal microvascular abnormalities, which can indicate the presence of Diabetic Retinopathy, or other pathologies. Last but not least, vessel segmentation can also help to remove false detections of lesions like microaneurysms, exudates and hemorrhages.

The vessels have a lower reflectance compared to other retinal structures, and thus they appear darker than the background in fundus images. The major vessels start vertically from the optic disc. The main difficulties of vessel segmentation come from the peripheral branches of the retinal vasculature. They have low contrast and sometimes are disconnected (Fig. B.1(a)). Other problems are due to the noise, which can be introduced by many reason, such as the reflections of other retinal layers or structures. It can disconnect the vessels or create vessel-like patterns (Fig. B.1(b),B.1(c),B.1(d)). It is very hard to segment all vessels. However, in order to identify retinal objects like the optic disc, only a gross main vascular network segmentation is needed. On the other hand, to eliminate false detections of lesions like microaneurysms, small branches and low contrast vessel is needed. The proposed methods are based on these principles.

B.1 State of the art

Numerous studies concerning vessel segmentation exist in the literature, which can be classified in several categories. Methods based on mathematical morphology are known as efficient and noise resistant, by identifying specific shapes. For instance, a rotating linear structuring element has been used to extract vessels (Spencer et al. [98], Walter and Klein [107]). Matched filtering is another widely used method (Chaudhuri et al. [16], Gang et al. [36], Hoover et al. [49]). Usually 2D matched filters, for example, a set of Gaussians, are used to match the profile of vessels. However, the kernels are often convolved with the image several times in different angles, which cause a computational overhead. The

APPENDIX B. VESSEL SEGMENTATION




Figure B.1: Examples of difficult cases for vessel segmentation. a) Low contrast and disconnected vessel, b) Bright noises interrupt the vessels, c) Elongated reflections create fake vessels, d) Choroid vessels.

standardized kernel may not have a good response for vessels of different sizes and profiles. Methods based on vessel tracking can provide more accurate segmentations (Can et al. [15], Chutatape et al. [18], Tolia and Panas [104]). Two steps are involved in such methods: starting points detection and tracking or growing method. Grid searching for local minima or matched Gaussian filters can provide initial points. Fuzzy C-means clustering on vessel profiles or Kalman filters have been used to track vessels. In addition, these methods can provide the width and bifurcation information which is very useful for extracting vessel characteristics. But problems can arise from disconnected vessels. Al-Diri et al. [6] used a similar method based on active contours. A generalized morphological order filter was used to identify vessels centerlines. Then an active contour model “Ribbon of Twins” was used to segment vessels. Supervised classification methods can use many powerful statistical analysis tools. Marín et al. [63] used a pixel-based neural network with 7 features based on gray level and moment invariants to segment vessels. Staal et al. [99] started by extracting image ridges using a Hessian

matrix. Then a k-nearest neighbor classifier was used to perform the classification. The problem of supervised methods is that a database with good manual vessels segmentation masks is needed, which is difficult to obtain for vessels.

Two vessel segmentation methods are proposed. According to the usage, firstly, a very efficient method based on the proposed pre-processing method (Chapter 2) is proposed to segment main vessels. Secondly, a method based on Gabor filters, inspired by the work of Guillaume Thibault, will be introduced, which is capable to segment and connect small low contrasted vessels.

B.2 Method based on alternating sequential filter

The objective of this method is to efficiently segment the main vessels of the retina. It is inspired by methods based on the residue of a morphological closing (see for instance Walter and Klein [107]), but here we use an alternate sequential filter (ASF, Sternberg [100] and Serra [93]), which has been introduced in Chapter 2.

The negative residue between the original image (green channel) and the result of the ASF keeps all dark structures with details, called pre-processed image. An example is shown in Fig. B.2. The size of the filter goes from 1 to d_{vessel} (the maximum width of the vessels, see Chapter 3).




Figure B.2: Vessels segmentation based on alternating sequential filter. a) Original green channel, b) Negative residues of alternating sequential filter. c) Result of automatic thresholding, d) Result of area opening.

Then, a threshold is applied on the pre-processed image. In order to automatically choose the threshold value, we choose it in such a way that the resulting mask represents around 13% of the field of view (FOV) area. However, we consider a minimum threshold value of 5, in order to take into account images with few or no vessels. The thresholded image is shown in Fig. B.2(c). An area opening (with a size criterion of d_{vessel}^2) removes isolated small structures. Finally, a closing with a hexagonal structuring element is used to fill small holes inside large vessels (Fig. B.2(d)). The entire image is shown in Fig. B.3.

This method is extremely efficient. The final vessel mask contains the main vessels. But small vessels and low contrast vessels are lost. This result provides sufficient information for retinal structures identification, for instance, the optic disc. The following section will introduce a method based on the Gabor filter to give a finer segmentation.




Figure B.3: Vessels segmentation based on alternating sequential filter. a) Original image (green channel), b) Segmentation result.

B.3 Method based on Gabor filter

The above presented efficient method gives a good segmentation on the main vessels. However, it does not give details on small vessel branches, especially isolated and low contrasted ones. A method based on Gabor filter is presented in this section. It solves these problems and is very useful for removing false detections of lesions. It is based on work by Guillaume Thibault.

B.3.1 Gabor filter

The Gabor filter was proposed by Gabor [35] and was initially used for signal processing. Daugman et al. [21] found that Gabor filters are similar to some processing steps of the human visual system, concerning frequency and orientation representations. A 2D Gabor filter is commonly used for edge detection and texture representation and discrimination. It has succeeded in many applications, such as texture analysis, iris, face, fingerprint and numerals recognition. A 2D Gabor filter is a Gaussian kernel function modulated by a sinusoidal plane wave in the spatial domain. Its real part is given by:

$$g(x,y; \lambda, \theta, \psi, \sigma, \gamma) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \cos\left(2\pi\frac{x'}{\lambda} + \psi\right) \quad (\text{B.1})$$

where $x' = x \cos \theta + y \sin \theta$ and $y' = -x \sin \theta + y \cos \theta$. λ is the wavelength, θ is the orientation, ψ is the phase, σ is the standard deviation of the Gaussian function and γ is the aspect ratio.

B.3.2 Vessels extraction

Due to the elliptical Gaussian shape of the Gabor filter, the vessels have high responses to the Gabor filter in the corresponding orientation. Thus, we propose to use 12 Gabor filters with different orien-




Figure B.4: Gabor Kernel at orientation $\frac{\pi}{4}$.

tations. We tune the coefficients of the Gabor filters as follows: $\lambda = d_{vessel}$, $\psi = 0$, $\sigma = d_{vessel}/12$, $\gamma = 2$ and $\theta = \frac{\pi}{12}N$, $N = \{0,1 \dots 11\}$. An example of the tuned kernels is shown in Fig. B.4.

The Gabor filters are applied to the pre-processed image (I_{pre}), as the previous proposed method. The I_{pre} of the same test image is shown in Fig. B.5(a). Fig. B.5(b) shows the I_{pre} in a heatmap look-up table for a better visualization. The results of Gabor filters at orientations 0 (horizontal) and $\frac{\pi}{2}$ (vertical) are shown in Fig. B.5(c) and Fig. B.5(d). Notice the underlined structures, a vertical (near left border of the image) and a horizontal vessel. They are enhanced and connected by the Gabor filter at its corresponding orientation.




Figure B.5: Results of Gabor filters. a) Details of I_{pre} , b) Heatmap of I_{pre} , c) Result of Gabor filter at horizontal orientation, d) Result of Gabor filter at vertical orientation.

In order to segment the vessels from the 12 filtered images, based on the fact that vessels are enhanced by the corresponding Gabor filter, the first possible approach is to take the supremum of the 12 filtered images (Fig. B.6(a)). However, while trying to threshold this image (for instance, Fig. B.6(b) shows the thresholding result at 22), the large vessels appear too wide, and contrarily, the same threshold is not low enough to segment small vessels. Moreover, threshold varies between different images.

We propose a method which uses the orientation information instead of the intensity information from the results of the Gabor filters. For each pixel p , the orientation α_p is defined by the orientation of the Gabor filter which has the maximum response:

$$\alpha_p = \underset{\alpha}{\operatorname{argmax}} \operatorname{Gabor}_{\alpha}(p) \tag{B.2}$$




Figure B.6: Thresholding on the supremum of the results of Gabor filters. a) Supremum image, b) Thresholding at 22.

An orientation map is generated (Fig. B.7(a)). Because the pixels inside the vessels have the maximum response at the same orientation along the vessel, they appear in the color map as a connected component with large area. If we merge the pixels with the adjacent orientations together, they will have a larger area and be more homogeneous. Thus, pixels of every three adjacent orientations are merged together. For example, Fig. B.7(b) shows the combined image of the horizontal orientation (merged with $\frac{\pi}{12}$ and $\frac{11\pi}{12}$). An area opening is used to remove objects smaller than $d_{vessel}^2/4$. And a closing by a hexagonal structuring element with size 2 to fill small holes inside the vessels (Fig. B.7(c)).

For each connected component of the remaining structures, several features are extracted and an empirical selection is done by a set of criteria. The features are listed as follows:

- Geodesic length (*geoLength*). In the meantime, the coordinates of the two endpoints are noted for other computations.
- Ratio between Euclidean distance and geodesic length (R_{dist}). This feature reflects the tortuosity of the structure.
- Mean width (*meanWidth*), max width (*maxWidth*) and width variance (*widthVar*) along the main axis of the structure. The main axis is defined by the segment between the endpoints obtained while computing the geodesic length. An illustration is shown in Fig. B.8(a). P_1 and P_2 are the endpoints of the geodesic arc (blue line). The coordinate system is transferred to the main axis (red line). For every point along the main axis, the width is computed.
- Mean intensity (*meanVessel*). It is computed on I_{pre} .
- Orientation difference (β_p). First, the orientation of the connected component is calculated from the two endpoints. Then the difference between it and the its corresponding α_p in the orientation map are computed and denoted β_p . In normal cases, this angle is nearly zero, because the analyzed structures are selected by their orientations in the beginning. But exceptions exist like the one pointed by the red arrow in Fig. B.7(c), which is orthogonal to α_p . This structure is adjacent to a vertical vessel and does not exist in the original image. It appears here because in the blank zones next to the vessel borders, the Gabor filter which has an orthogonal orientation to the vessel and touches the vessel, always has the maximum response. Fig. B.8(b) illustrates




Figure B.7: Orientation information from Gabor filters. a) Orientation map of each pixel, b) Selected pixels at orientation $\frac{11\pi}{12}$, 0 and $\frac{\pi}{12}$, c) Area opening removes small objects and closing fills holes inside vessels, d) Selection by features. e) Union of the selection results at 12 orientations and . f) Original green channel

the forming of a vertical zone of horizontal α_p against a vertical vessel. In this case, β_p reaches its maximum value $\frac{\pi}{2}$.

Then a decision tree is applied with empirical criteria. The structures satisfying the following conditions are kept as vessels:

- $geoLength > \frac{1}{2}d_{FOV}$
- $\beta_p \leq 0.4$ and $meanVessel \geq 3$ and $maxWidth < 2d_{vessel}$
- $\beta_p \leq 0.4$ and $meanVessel \geq 3$ and $widthVar < 0.06$ and $\frac{maxWidth}{meanWidth} < 2.5$
- $\beta_p \leq 0.4$ and $meanVessel > 1$ and $\frac{1}{2}maxWidth < d_{vessel}$ and $R_{dist} > 0.85$ and $geoLength > 2d_{vessel}$

The selection result on Fig. B.7(c) is shown in Fig. B.7(d). The same operation is repeated for the 12 orientations. Then the results are united together to become the vessel mask (Fig. B.7(e)). By comparing to the original image, tiny vessels and low contrasted vessels are segmented. And their width seems to be nicely correlated to the width of the vessels. However, the high sensitivity of this method can also cause false detections, for example, the three structures pointed by red arrows in Fig. B.7(e) are not vessels. They are caused by other dark structures. The left one is actually

a microaneurysm. We keep all of them in the vessel mask here, as we mentioned that this vessels segmentation will be used to remove false lesion detections afterwards. The methods to distinguish the false vessels segmentation and the real lesions will be detailed in each chapter of lesions detection. The general idea is to use the intensity variances within these structures.




Figure B.8: a) Illustration of width features extraction. b) An example of a vertical zone (blue) of horizontal α_p .

The final segmentation result is illustrated in Fig. B.9(c). The result of the ASF based method is also shown for comparison. The result of Gabor filter based method contains much more details where the vessels are small or low contrasted. Moreover, the vessel mask has smoother borders and corresponds exactly to the real width. Thus the thresholding problems are, in a manner, avoided. On the other hand, this method is too sensitive compared to the other methods. For example, in the macula region, there are false segmentations.

B.4 Segmentation results on bad quality images

We choose another two “difficult” images in our database to test the two proposed methods. Fig. B.10(a) is the green channel of the first image. It has a bright and low contrast zone in the right, and contains exudates and hemorrhages. The result of the ASF contains large vessels without much details. The false segmentations in the red window (Fig. B.10(b)) come from hemorrhages. The result of the Gabor filter based method contains more details, for example, the small vessels pointed by red arrow in Fig. B.10(c), which has very low contrast. But it is also sensitive to other dark structures, for example, in the same zone inside red window.

Another difficult case is shown in Fig. B.11(a) with a bright zone in the top and a dark zone in the bottom. In addition, high contrasted exudates and hemorrhages appear together in the center left of the image. The result of the ASF keeps a few large vessels with negligible noise. The result of the Gabor filters gives more details, but also more noise.

As vessel segmentation methods, there are advantages and disadvantages to both approaches. But the disadvantages may not be important or can be eliminated for certain usages of the vessel masks.




Figure B.9: Vessels segmentation results comparison. a) Original image, b) Segmentation by ASF, c) Segmentation by Gabor filters.




Figure B.10: Vessels segmentation results comparison. a) Original image, b) Segmentation by ASF, c) Segmentation by Gabor filters.




Figure B.11: Vessels segmentation results comparison. a) Original image, b) Segmentation by ASF, c) Segmentation by Gabor filters.

Optic Disc localization

The Optic Disc (OD) is the entry point of blood vessels. In color eye fundus images, it appears as a bright round region. The OD localization provides anatomic information and is the basis for some specific diseases detection. It can also help to remove false detections of lesions, like exudates. Although the OD is the biggest and brightest structure in fundus images, its contour is somehow ambiguous. It can be hidden by major blood vessels or fade out in the background. We concentrate in the localization of OD in this work by identifying the center of the disc. The main objective is to help removing false lesions detections. In real conditions, the OD could appear in any position in the retinal image or even be absent from the image. In some cases, it is not well contrasted and not the brightest region in the image (Fig. C.1).

C.1 State of the art

A number of OD detection methods are based on the assumption that the OD is the brightest structure in the image, for example, Walter and Klein [107]. After a local contrast enhancement, the brightest and largest connected component could be considered as the OD. However, several images in our database do not satisfy this condition. There are also methods based on the round shape of the OD. The Hough transform is the most used method (Abdel-Ghafar et al. [1]). But again, in our case, the optic disc is often partially hidden. Another family of methods is based on the fact that the vessels enter and leave the retina through the OD. Methods based on vessel model fitting or vessel template filtering are used in Abdel-Razik Youssif et al. [2], Chutatape [17], Foracchia et al. [32]. Lam and Chutatape [58] proposed a vessel tracking method to find the OD. Based on the fact that the OD is usually located close to large vessels, Tobin et al. [103] use the vessel density, thickness and orientation information. Similarly, Hoover and Goldbaum [48] proposed a voting system by fuzzy convergence of the blood vessels. Besides, there also exist other methods based on mathematical morphology (Hajer et al. [43]) and template matching (Lalonde et al. [57]). After testing these different ideas on our database, we have finally found that the combination of the intensity information with the position of the main vessels leads to a good solution for our application.




Figure C.1: Examples of difficult cases for OD detection.

C.2 Border reflection detection

Reflections in the border of the field of view are very common in fundus images (Fig. C.2(a)). They could cause false OD detection. We choose to detect them separately in the beginning. They are clear in green channel, as other structures and lesions. But we find that they are also visible in the blue channel (Fig. C.2(b)). Because the blue channel contains less information on other structures, it is good for reflection detection. Here we use the average of the green and blue channels as the input (Fig. C.2(c)).

A mean filter is applied to the original image with window size $d_{FOV}/7$ to get a background estimation. Then, the resulting image is subtracted from the original image (Fig. C.3(a)). The outline of the field of view is taken as the marker to do a reconstruction under the previous image. The bright structures in the central region are removed (Fig. C.3(b)). Then a thresholding at 3 gives the binary mask containing border reflections (Fig. C.3(c)). A length opening is used to remove small and short structures. The length criterion is set to $d_{FOV}/4$ (Fig. C.3(d)).

The final step ensures that the segmented reflections are in the border region. We define a central region by eroding the mask of the field of view by a hexagonal structuring element of size $d_{FOV}/10$ (illustrated by a red circle in Fig. C.3(d)). The image is divided into two parts: inner region (inside the eroded mask) and outer region (outside the eroded mask). Then, the number of the pixels of the




Figure C.2: Border reflection. a) Green channel, b) Blue channel, c) Average of green and blue channel.




Figure C.3: Detection of border reflections. a) Background removal image, b) Reconstruction image from the border circle, c) Thresholded image, d) Result of opening by length.

first segmentation result is counted in both regions. A ratio of the inner area over the outer area is computed. If it is larger than 0.25, a reconstruction is done from the eroded FOV mask (inner region) to remove the structures touching it. If the ratio is less than 0.25, the structures inside the central region will also be removed directly, but the outside structures are kept.

C.3 Intensity and size information

Firstly, we use the intensity and size information to get OD candidates. The average of the three original color channels is taken as the input image (I_{input} , Fig. C.4(a)). A closing is applied to remove dark structures like vessels (Fig. C.4(b)). Then, a large mean filter with kernel size $d_{FOV}/7$ is used to estimate the background, and is subtracted from the previous image (Fig. C.4(c)). After that, the detected border reflection is subtracted from this image (Fig. C.4(d)). Another mean filter with size d_{OD} is applied to smooth the resulting image (Fig. C.4(e)). Then a thresholding starts from the maximum intensity value of this image and is lowered until the area of the thresholded image reaches $0.6 \times d_{OD}^2$. A binary mask of the OD candidates is obtained, denoted $I_{ODcandi}$ (Fig. C.4(f)). In this example, we have only one candidate, but it is possible to have several.




Figure C.4: Intensity analysis. a) Input image, the average of three channels, b) Dark structures removal image, c) Background removal image, d) Border reflection removal image, e) Mean filtered image, d) Thresholded image.

C.4 Vessel analysis

The segmented vessels will be used to select and verify the OD candidates. Any state of the art vessels segmentation method will be fine. Here we use the method based on the alternate sequential filter (see Appendix B) for the sake of efficiency. Besides, only main vessels are needed here. An image example and its vessels mask are shown in Fig. C.5(a) and Fig. C.5(b).

From the vessels mask, width and orientation information are extracted based on a local skeleton analysis (Fig. C.5(c)). The skeleton algorithm from Vincent [106] is implemented. Again, different state of the art method would probably work here. Based the fact that vessels usually extend vertically after coming out of the OD, we first compute the vertical projection of the vertical vessels whose width is between $0.4d_2$ and d_2 . The profile of such a projection is shown in Fig. C.6. The position of its maximal response (x_{proj}) gives an estimation of the OD horizontal position. Based on the fact that vessel density is large in the OD region, a rectangular window (Fig. C.5(b)) is used to compute a vessel




Figure C.5: Skeleton of vessels mask. a) Original image, b) Segmented vessel mask, c) Skeleton.




Figure C.6: Vessel projection to the horizontal axis.

density map, as proposed by Tobin et al. [103]. The location which has the largest density gives a second estimation of the OD position (x_{dens}, y_{dens}) .

The ratio between the number of the vertical vessel pixels (N_v) and two other diagonal directions (N_{d1}, N_{d2}) is computed as an important feature, denoted R_V . It is very useful for the case where the OD is absent. For example, Fig. C.7 shows an image without OD and the OD candidate obtained by intensity analysis. We find that in such images, there are few vessels, especially vertical vessels. In order to illustrate this point, we have computed these values for a normal image (Fig. C.4(a)) and an image without OD (Fig. C.7(a)). The result is summarized in Tab. C.1. The total number of pixels is sharply reduced in the image without OD, especially for vertical vessels.




Figure C.7: Image without OD. a) Average of three channels, b) Candidate from the result of intensity analysis.

	N_v (vertical vessel)	N_{d1}	N_{d2}
Normal image (Fig. C.4(a))	2498	2860	3443
Image without OD (Fig. C.7(a))	281	1402	1745

Table C.1: Number of pixels in skeleton image in three directions.

C.5 Selection

The final selection is performed based on $I_{ODcandi}$, x_{proj} , x_{dens} and R_V . First of all, we cope with images from which the OD is absent. If $N_v < 300$ **or** $R_V \leq 0.2$, we conclude that there is no OD in the image and a void region is returned. Otherwise, we consider two situations:

- If $|x_{proj} - x_{dens}| < d_{OD}$, we take $x_{OD} = (x_{proj} + x_{dens})/2$. Then, we search for the vertical position of the OD with a moving window ($d_{OD} \times d_{OD}$ pixels), along the y-axis on I_{input} , horizontally limited to $[x_{OD} - 0.5d_{OD}, x_{OD} + 0.5d_{OD}]$. The y-coordinate of the position with the maximum mean value within the window is noted y_{OD} . This process is illustrated in Fig. C.8. If there is a candidate OD within $0.5d_{OD}$ pixels of (x_{OD}, y_{OD}) , it is considered as the detected OD. Otherwise, a void region is returned.
- If $|x_{proj} - x_{dens}| \geq d_{OD}$, then we look for a convenient OD candidate around position x_{proj} , using the method described above. If no candidate is found, then the procedure is repeated, starting this time with x_{dens} . Finally, if no convenient candidate is found, a void region is returned.

C.6 Evaluation and conclusions

We evaluated the proposed algorithm on a sub-set of e-optha, which contains 1044 images with various resolutions and qualities. The OD detection algorithm runs automatically on the entire base and gives the coordinates of the OD location. The results are compared to a manual annotation. If the distance between them is less than $d_{OD}/3$, the OD is well localized. In 964 images, the OD is correctly localized,




Figure C.8: Estimation of OD position using vessel information.

which achieves an sensitivity of 92.3%. Among 1044 images, there are 10 images without the presence of OD. They are all marked as no OD. Thus, a specificity of 100% is achieved. Among the 80 images where the algorithm fails to detect the OD, the vessel information gives a right prediction of the OD position on 10 images. Most missed detections are caused by the low contrast of the OD and the presence of other bright structures in the image (for example, Fig. C.9).



Figure C.9: An image on which the algorithm fails to detect the OD.

In this section, we proposed an OD localization method. It combines different features, such as the intensity, size, and vessel information. The result of vessel analysis performs a selection and a confirmation on the OD candidates obtained by intensity and size analysis. This method is capable to decide whether there is an OD or not presented in the image with a high accuracy. Because of the use of spatial calibration, the image can be adjusted to arbitrary size, without changing anything of the algorithm. Thus, images can be resized to a small scale while extracting OD candidates.

The majority of errors are due to the low contrast of the OD. Note that the OD detection, in this work, is used to remove the false detections of the lesions. However, if the OD is not detected because of its low contrast, it has a small chance to be taken as a bright lesion candidate. Moreover, failure to detect an OD is less harmful than take a lesion, for example a large exudate, as an OD.

APPENDIX C. OPTIC DISC LOCALIZATION

Future work should focus on the cases where the OD has a low contrast. If the vascular network is well segmented, the OD localization is totally possible depending only on the vessel information. Actually, we tried to fit a vessel model to estimate the OD location. But it largely increases the computational work, with little improvement. Other vessel features like symmetry and similarity could also be used. Texture features is another possibility. Even if its contrast is very low, the OD region tends to have a smoother texture, and a higher gradient along its edge. The round shape, of course, is another useful feature. But it is always cut by vessels or only partially visible.

Another perspective is the localization of the macula. It's more difficult to localize the macula than the OD. First of all, its contrast is usually very low, and its contour is not clearly defined. Moreover, it is quite often absent from the image. Once we detect the OD, it helps to localize the macula, because the distance between them is known. We proposed a method based on this fact. First the same method for OD candidates extraction is applied for macula detection. Then, the definitive distance to the detected OD is used to select the candidate. This method works on good quality images. It will be used in practice for removing false detection of hemorrhages. Suppose the macula is detected, further work such as defining left or right eye, is easily achieved. However, this is a chicken-egg problem. If the left or right eye is defined beforehand either manually or automatically, it helps to localize the macula. In fact, there are works taking this information as priori input of their algorithms.

ASF

alternating sequential filter. 23

AUC

area under curve. 57

d_{FOV}

Field-of-view width. 34, 62

diabetic retinopathy

Complication of diabetes which causes damage to the retina. 11

d_{MA}

Maximum width of microaneurysms. 34, 62

d_{OD}

Optic disc height. 34

d_{vessel}

Maximum vessel width. 34

FN

false negative. 55, 77, 98

FOV

field of view. 31, 62

FP

false positive. 55, 77, 99

γ

Opening. 22

γ_B^n

Opening of size n with structuring element B. 40

γ^{rec}

Opening by reconstruction. 23

OPHDIAT

A telemedicine network for diabetic retinopathy screening established in 2004 by Assistance Publique - Hôpitaux de Paris (AP-HP). 11

φ

Closing. 22

φ_B^n

Closing of size n with structuring element B. 40

φ^{rec}

Closing by reconstruction. 23

PPV

positive prediction value. 56

$Rec_X(Y)$

Reconstruction under X by taking Y as marker. 41

ρ

Top-Hat extension: $\rho(I) = I - \{\gamma\varphi(I) \wedge I\}$, based on an ASF. 22

ρ_{asf}

Top-Hat extension: $\rho_{asf}(I) = I - \{\varphi_i\gamma_i \cdots \varphi_2\gamma_2\varphi_1\gamma_1\varphi(I) \wedge I\}$, based on an ASF. 23

ROC

receiver operating characteristic. 56, 80, 99

τ

Top-Hat: $\tau(I) = I - \gamma(I)$. 22

TeleOphta

A collaborative research project, funded by the National Agency of Research (ANR) through the TecSan program. Its main objective is to develop the tools for fundus image classification, in order to assist in the detection of retinal diseases, especially of diabetic retinopathy. 11

TN

true negative. 56, 77

TP

true positive. 55, 77, 97

- [1] R. Abdel-Ghafar, T. Morris, T. Ritchings, and I. Wood. Detection and characterisation of the optic disk in glaucoma and diabetic retinopathy. In *Proceedings of Medical Image Understanding and Analysis*, 2004. [129](#)
- [2] A. Abdel-Razik Youssif, A. Ghalwash, and A. Abdel-Rahman Ghoneim. Optic disc detection from normalized digital fundus images by means of a vessels' direction matched filter. *Medical Imaging, IEEE Transactions on*, 27(1):11–18, 2008. [129](#)
- [3] S. Abdelazeem. Micro-aneurysm detection using vessels removal and circular hough transform. In *Radio Science Conference, 2002.(NRSC 2002). Proceedings of the Nineteenth National*, pages 421–426. IEEE, 2002. [63](#)
- [4] M. D. Abràmoff, M. Niemeijer, M. S. A. Suttorp-Schulten, M. A. Viergever, S. R. Russell, and B. v. Ginneken. Evaluation of a system for automatic detection of diabetic retinopathy from color fundus photographs in a large population of patients with diabetes. *Diabetes Care*, 31(2): 193–198, Jan. 2008. ISSN 0149-5992, 1935-5548. doi: 10.2337/dc07-1312. [10](#)
- [5] C. Agurto, E. S. Barriga, V. Murray, S. Nemeth, R. Crammer, W. Bauman, G. Zamora, M. S. Pattichis, and P. Soliz. Automatic detection of diabetic retinopathy and age-related macular degeneration in digital fundus images. *Investigative Ophthalmology & Visual Science*, 52(8): 5862–5871, Jan. 2011. ISSN , 1552-5783. doi: 10.1167/iovs.10-7075. [10](#)
- [6] B. Al-Diri, A. Hunter, and D. Steel. An active contour model for segmenting and measuring retinal vessels. *Medical Imaging, IEEE Transactions on*, 28(9):1488–1497, 2009. [118](#)
- [7] B. Antal and A. Hajdu. An ensemble-based system for microaneurysm detection and diabetic retinopathy grading. *Biomedical Engineering, IEEE Transactions on*, 59(6):1720–1726, 2012. [63](#), [77](#)
- [8] B. Antal and A. Hajdu. Improving microaneurysm detection in color fundus images by using context-aware approaches. *Computerized Medical Imaging and Graphics*, 2013. [18](#), [63](#)
- [9] A. Aquino Martín. *Fundus digital image processing: automated segmentation of the main retinal anatomical structures*. PhD thesis, 2011. [19](#)

BIBLIOGRAPHY

- [10] J. P. Bae, K. G. Kim, H. C. Kang, C. B. Jeong, K. H. Park, and J.-M. Hwang. A study on hemorrhage detection using hybrid method in fundus images. *Journal of digital imaging*, 24(3): 394–404, 2011. [84](#)
- [11] S. Beucher. Numerical residues. *Mathematical Morphology: 40 Years On*, pages 23–32, 2005. [19](#), [24](#), [47](#)
- [12] A. Bhalerao, A. Patanaik, S. Anand, and P. Saravanan. Robust detection of microaneurysms for sight threatening retinopathy screening. In *Computer Vision, Graphics & Image Processing, 2008. ICVGIP'08. Sixth Indian Conference on*, pages 520–527. IEEE, 2008. [63](#)
- [13] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001. [27](#), [43](#), [50](#), [73](#), [94](#)
- [14] P. C. Bunch, J. F. Hamilton, G. K. Sanderson, and A. H. Simmons. A free response approach to the measurement and characterization of radiographic observer performance. In *Application of Optical Instrumentation in Medicine VI*, pages 124–135. International Society for Optics and Photonics, 1977. [77](#), [95](#)
- [15] A. Can, H. Shen, J. N. Turner, H. L. Tanenbaum, and B. Roysam. Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms. *Information Technology in Biomedicine, IEEE Transactions on*, 3(2):125–138, 1999. [118](#)
- [16] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson, and M. Goldbaum. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Transactions on medical imaging*, 8(3):263–269, 1989. [117](#)
- [17] O. Chutatape. Fundus foveal localization based on vessel model. In *Engineering in Medicine and Biology Society, 2006. EMBS'06. 28th Annual International Conference of the IEEE*, pages 4440–4444. IEEE, 2006. [129](#)
- [18] O. Chutatape, L. Zheng, and S. Krishnan. Retinal blood vessel detection and tracking by matched gaussian and kalman filters. In *Engineering in Medicine and Biology Society, 1998. Proceedings of the 20th Annual International Conference of the IEEE*, volume 6, pages 3144–3149. IEEE, 1998. [118](#)
- [19] M. J. Cree, J. A. Olson, K. C. McHardy, P. F. Sharp, and J. V. Forrester. A fully automated comparative microaneurysm digital detection system. *Eye*, 11(5):622–628, 1997. [62](#)
- [20] G. Danaei, M. M. Finucane, Y. Lu, G. M. Singh, M. J. Cowan, C. J. Paciorek, J. K. Lin, F. Farzadfar, Y.-H. Khang, G. A. Stevens, et al. National, regional, and global trends in fasting plasma glucose and diabetes prevalence since 1980: systematic analysis of health examination surveys and epidemiological studies with 370 country-years and 2·7 million participants. *The Lancet*, 378(9785):31–40, 2011. [9](#)
- [21] J. G. Daugman et al. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Optical Society of America, Journal, A: Optics and Image Science*, 2(7):1160–1169, 1985. [120](#)
- [22] E. Decencière, G. Cazuguel, X. Zhang, G. Thibault, J.-C. Klein, F. Meyer, B. Marcotegui, G. Quellec, M. Lamard, R. Danno, et al. Teleophta: Machine learning and image processing methods for teleophthalmology. *IRBM*, 2013. [9](#), [10](#), [108](#)

- [23] C. Delcourt, P. Massin, and M. Rosilio. Epidemiology of diabetic retinopathy: Expected vs reported prevalence of cases in the french population. *Diabetes & metabolism*, 35(6):431–438, 2009. 9
- [24] B. Dupas, T. Walter, A. Erginay, R. Ordonez, N. Deb-Joardar, P. Gain, J.-C. Klein, and P. Massin. Evaluation of automated fundus photograph analysis algorithms for detecting microaneurysms, haemorrhages and exudates, and of a computer-assisted diagnostic system for grading diabetic retinopathy. *Diabetes & metabolism*, 36(3):213–220, 2010. 10, 30, 52
- [25] A. Erginay, A. Chabouis, C. Viens-Bitker, N. Robert, A. Lecleire-Collet, and P. Massin. Ophdiat: Quality-assurance programme plan and performance of the network. *Diabetes & metabolism*, 34(3):235–242, 2008. 9
- [26] P. Feng, Y. Pan, B. Wei, W. Jin, and D. Mi. Enhancing retinal image by the contourlet transform. *Pattern Recognition Letters*, 28(4):516–522, 2007. 17
- [27] N. I. Fisher. *Statistical analysis of circular data*. Cambridge University Press, 1995. 88
- [28] A. Fleming, S. Philip, K. Goatman, G. Williams, J. Olson, and P. Sharp. Automated detection of exudates for diabetic retinopathy screening. *Physics in medicine and biology*, 52:7385, 2007. 37
- [29] A. Fleming, K. Goatman, G. Williams, S. Philip, P. Sharp, and J. Olson. Automated detection of blot haemorrhages as a sign of referable diabetic retinopathy. In *Proc. Medical Image Understanding and Analysis*, 2008. 30, 84, 92
- [30] A. D. Fleming, S. Philip, K. A. Goatman, J. A. Olson, and P. F. Sharp. Automated microaneurysm detection using local contrast normalization and local vessel detection. *Medical Imaging, IEEE Transactions on*, 25(9):1223–1232, 2006. 30, 63, 70
- [31] A. D. Fleming, S. Philip, K. A. Goatman, J. A. Olson, and P. F. Sharp. Automated assessment of diabetic retinal image quality based on clarity and field definition. *Investigative ophthalmology & visual science*, 47(3):1120–1125, 2006. 117
- [32] M. Foracchia, E. Grisan, and A. Ruggeri. Detection of optic disc in retinal images by means of a geometrical model of vessel structure. *Medical Imaging, IEEE Transactions on*, 23(10):1189–1195, 2004. 129
- [33] M. Foracchia, E. Grisan, and A. Ruggeri. Luminosity and contrast normalization in retinal images. *Medical Image Analysis*, 9(3):179–190, 2005. 17
- [34] A. J. Frame, P. E. Undrill, M. J. Cree, J. A. Olson, K. C. McHardy, P. F. Sharp, and J. V. Forrester. A comparison of computer based classification methods applied to the detection of microaneurysms in ophthalmic fluorescein angiograms. *Computers in Biology and Medicine*, 28(3):225–238, 1998. 62
- [35] D. Gabor. Theory of communication. part 1: The analysis of information. *Electrical Engineers-Part III: Radio and Communication Engineering, Journal of the Institution of*, 93(26):429–441, 1946. 120

BIBLIOGRAPHY

- [36] L. Gang, O. Chutatape, and S. M. Krishnan. Detection and measurement of retinal vessels in fundus images using amplitude modified second-order gaussian filter. *Biomedical Engineering, IEEE Transactions on*, 49(2):168–172, 2002. [117](#)
- [37] M. García, M. I. López, D. Álvarez, and R. Hornero. Assessment of four neural network based classifiers to automatically detect red lesions in retinal images. *Medical engineering & physics*, 32(10):1085–1093, 2010. [84](#)
- [38] G. Gardner, D. Keating, T. Williamson, and A. Elliott. Automatic detection of diabetic retinopathy using an artificial neural network: a screening tool. *British Journal of Ophthalmology*, 80(11):940–944, 1996. [37](#)
- [39] L. Giancardo, F. Mériaudeau, T. P. Karnowski, K. W. Tobin, Y. Li, and E. Chaum. Microaneurysms detection with the radon cliff operator in retinal fundus images. In *SPIE Medical Imaging*, pages 76230U–76230U. International Society for Optics and Photonics, 2010. [63](#)
- [40] L. Giancardo, F. Meriaudeau, T. Karnowski, Y. Li, K. Tobin, and E. Chaum. Automatic retina exudates segmentation without a manually labelled training set. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 1396–1400. IEEE, 2011. [37](#), [53](#), [54](#), [55](#), [56](#)
- [41] L. Giancardo, F. Meriaudeau, T. Karnowski, Y. Li, S. Garg, K. Tobin, and E. Chaum. Exudate-based diabetic macular edema detection in fundus images using publicly available datasets. *Medical Image Analysis*, 16(1):216–226, 2012. [37](#), [52](#), [56](#), [57](#)
- [42] R. C. Gonzalez, R. E. Woods, and S. L. Eddins. *Digital image processing using MATLAB*, volume 2. Gatesmark Publishing Knoxville, 2009. [18](#)
- [43] J. Hajer, H. Kamel, and E. Noureddine. Localization of the optic disk in retinal image using the watersnake. In *Computer and Communication Engineering, 2008. ICCCE 2008. International Conference on*, pages 947–951. IEEE, 2008. [129](#)
- [44] B. Harangi, B. Antal, and A. Hajdu. Automatic exudate detection with improved naïve-bayes classifier. In *Computer-Based Medical Systems (CBMS), 2012 25th International Symposium on*, pages 1–4. IEEE, 2012. [37](#)
- [45] Y. Hatanaka, T. Nakagawa, Y. Hayashi, M. Kakogawa, A. Sawada, K. Kawase, T. Hara, and H. Fujita. Improvement of automatic hemorrhages detection methods using brightness correction on fundus images. In *Proc. SPIE*, volume 6915, pages 69153E–1, 2008. [84](#)
- [46] H. J. A. M. Heijmans. Theoretical aspects of gray-level morphology. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 13(6):568–582, 1991. [24](#)
- [47] J. Hipwell, F. Strachan, J. Olson, K. McHardy, P. Sharp, and J. Forrester. Automated detection of microaneurysms in digital red-free photographs: a diabetic retinopathy screening tool. *Diabetic medicine*, 17(8):588–594, 2000. [30](#)
- [48] A. Hoover and M. Goldbaum. Locating the optic nerve in a retinal image using the fuzzy convergence of the blood vessels. *Medical Imaging, IEEE Transactions on*, 22(8):951–958, 2003. [17](#), [129](#)

- [49] A. Hoover, V. Kouznetsova, and M. Goldbaum. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *Medical Imaging, IEEE Transactions on*, 19(3):203–210, 2000. 117
- [50] R. Jones. Connected filtering and segmentation using component trees. *Computer Vision and Image Understanding*, 75(3):215–228, 1999. 24
- [51] G. D. Joshi and J. Sivaswamy. Colour retinal image enhancement based on domain knowledge. In *Computer Vision, Graphics & Image Processing, 2008. ICVGIP'08. Sixth Indian Conference on*, pages 591–598. IEEE, 2008. 17
- [52] V. S. Joshi, R. J. Maude, J. M. Reinhardt, L. Tang, M. K. Garvin, A. A. Sayeed, A. Ghose, M. U. Hassan, and M. D. Abramoff. Automated detection of malarial retinopathy-associated retinal hemorrhages. *Investigative Ophthalmology & Visual Science*, 53(10):6582–6588, 2012. 84
- [53] G. B. Kande, T. S. Savithri, P. V. Subbaiah, and M. Tagore. Detection of red lesions in digital fundus images. In *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on*, pages 558–561. IEEE, 2009. 83
- [54] G. B. Kande, T. S. Savithri, and P. V. Subbaiah. Automatic detection of microaneurysms and hemorrhages in digital fundus images. *Journal of digital imaging*, 23(4):430–437, 2010. 63
- [55] T. Kauppi, V. Kalesnykiene, J. Kamarainen, L. Lensu, I. Sorri, A. Raninen, R. Voutilainen, H. Uusitalo, H. Kälviäinen, and J. Pietilä. Diaretdb1 diabetic retinopathy database and evaluation protocol. *Proc. Medical Image Understanding and Analysis (MIUA)*, pages 61–65, 2007. 52
- [56] A. Khotanzad and Y. H. Hong. Invariant image recognition by zernike moments. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(5):489–497, 1990. 101
- [57] M. Lalonde, M. Beaulieu, and L. Gagnon. Fast and robust optic disc detection using pyramidal decomposition and hausdorff-based template matching. *Medical Imaging, IEEE Transactions on*, 20(11):1193–1200, 2001. 129
- [58] H. K. Lam and O. Chutatape. Blood vessel tracking technique for optic nerve localisation for field 1-3 color fundus images. In *Information, Communications and Signal Processing, 2003 and Fourth Pacific Rim Conference on Multimedia. Proceedings of the 2003 Joint Conference of the Fourth International Conference on*, volume 3, pages 1437–1441. IEEE, 2003. 129
- [59] C. Lantuéjoul and F. Maisonneuve. Geodesic methods in quantitative image analysis. *Pattern Recognition*, 17(2):177–187, 1984. 67
- [60] B. Laÿ. Analyse automatique des images angiofluorographiques au cours de la rétinopathie diabétique. *Ecole Nationale Supérieure des Mines de Paris, Centre de Morphologie Mathématique, Paris, France*, 1983. 62
- [61] I. Lazar and A. Hajdu. Microaneurysm detection in retinal images using a rotating cross-section based model. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 1405–1409. IEEE, 2011. 63

BIBLIOGRAPHY

- [62] M. W. MacArthur and J. M. Thornton. Conformational analysis of protein structures derived from nmr data. *Proteins: Structure, Function, and Bioinformatics*, 17(3):232–251, 1993. [88](#)
- [63] D. Marín, A. Aquino, M. E. Gegúndez-Arias, and J. M. Bravo. A new supervised method for blood vessel segmentation in retinal images by using gray-level and moment invariants-based features. *Medical Imaging, IEEE Transactions on*, 30(1):146–158, 2011. [118](#)
- [64] C. Marino, E. Ares, M. Penedo, M. Ortega, N. Barreira, and F. Gomez-Ulla. Automated three stage red lesions detection in digital color fundus images. *WSEAS Transactions on Computers*, 7(4):207–215, 2008. [63](#)
- [65] P. Massin and A. Erginay. *Rétinopathie diabétique*. Elsevier Masson, 2010. [9](#), [13](#)
- [66] P. Massin, A. Erginay, A. Gaudric, and E. scientifiques et médicales Elsevier. *Rétinopathie diabétique*. Ed. scientifiques et médicales Elsevier, 2000. [11](#), [13](#)
- [67] P. Massin, A. Chabouis, A. Erginay, C. Viens-Bitker, A. Lecleire-Collet, T. Meas, P.-J. Guillausseau, G. Choupot, B. André, and P. Denormandie. Ophdiat: A telemedical network screening system for diabetic retinopathy in the île-de-france. *Diabetes & metabolism*, 34(3):227–234, 2008. [9](#), [10](#)
- [68] F. Meyer and J. Serra. Contrasts and activity lattice. *Signal Processing*, 16(4):303–317, 1989. [114](#)
- [69] A. Mizutani, C. Muramatsu, Y. Hatanaka, S. Suemori, T. Hara, and H. Fujita. Automated microaneurysm detection method based on double ring filter in retinal fundus images. In *SPIE Medical Imaging*, pages 72601N–72601N. International Society for Optics and Photonics, 2009. [63](#)
- [70] V. Morard, E. Decencière, and P. Dokladal. Geodesic attributes thinnings and thickenings. *Mathematical Morphology and Its Applications to Image and Signal Processing*, pages 200–211, 2011. [46](#), [67](#)
- [71] M. Niemeijer, B. van Ginneken, J. Staal, M. S. Suttorp-Schulten, and M. D. Abràmoff. Automatic detection of red lesions in digital color fundus photographs. *Medical Imaging, IEEE Transactions on*, 24(5):584–592, 2005. [63](#)
- [72] M. Niemeijer, B. van Ginneken, S. Russell, M. Suttorp-Schulten, and M. Abràmoff. Automated detection and differentiation of drusen, exudates, and cotton-wool spots in digital color fundus photographs for diabetic retinopathy diagnosis. *Investigative ophthalmology & visual science*, 48(5):2260–2267, 2007. [37](#)
- [73] M. Niemeijer, B. Van Ginneken, M. J. Cree, A. Mizutani, G. Quéllec, C. I. Sánchez, B. Zhang, R. Hornero, M. Lamard, C. Muramatsu, et al. Retinopathy online challenge: automatic detection of microaneurysms in digital color fundus photographs. *Medical Imaging, IEEE Transactions on*, 29(1):185–195, 2010. [77](#)
- [74] G. K. Ouzounis, M. Pesaresi, and P. Soille. Differential area profiles: decomposition properties and efficient computation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(8):1533–1548, 2012. [24](#)

- [75] M. Pesaresi and J. A. Benediktsson. A new approach for the morphological segmentation of high-resolution satellite imagery. *Geoscience and Remote Sensing, IEEE Transactions on*, 39(2): 309–320, 2001. [24](#)
- [76] S. Philip, A. D. Fleming, K. A. Goatman, S. Fonseca, P. Mcnamee, G. S. Scotland, G. J. Prescott, P. F. Sharp, and J. A. Olson. The efficacy of automated ‘disease/no disease’ grading for diabetic retinopathy in a systematic screening programme. *British Journal of Ophthalmology*, 91(11):1512–1517, Jan. 2007. ISSN , 1468-2079. doi: 10.1136/bjo.2007.119453. [10](#)
- [77] S. Pradhan, S. Balasubramanian, and V. Chandrasekaran. An integrated approach using automatic seed generation and hybrid classification for the detection of red lesions in digital fundus images. In *Computer and Information Technology Workshops, 2008. CIT Workshops 2008. IEEE 8th International Conference on*, pages 462–467. IEEE, 2008. [84](#)
- [78] G. Quellec, M. Lamard, P. M. Josselin, G. Cazuguel, B. Cochener, and C. Roux. Optimal wavelet transform for the detection of microaneurysms in retina photographs. *Medical Imaging, IEEE Transactions on*, 27(9):1230–1241, 2008. [63](#)
- [79] G. Quellec, M. Lamard, G. Cazuguel, B. Cochener, and C. Roux. Adaptive nonseparable wavelet transform via lifting and its application to content-based image retrieval. *Image Processing, IEEE Transactions on*, 19(1):25–35, 2010. [101](#)
- [80] G. Quellec, M. Lamard, M. D. Abràmoff, E. Decencièrre, B. Lay, A. Erginay, B. Cochener, and G. Cazuguel. A multiple-instance learning framework for diabetic retinopathy screening. *Medical Image Analysis*, 2012. [101](#)
- [81] G. Quellec, M. Lamard, B. Cochener, Z. Droueche, B. Lay, A. Chabouis, C. Roux, and G. Cazuguel. Studying disagreements among retinal experts through image analysis. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pages 5959–5962. IEEE, 2012. [102](#)
- [82] G. Quellec, M. Laniard, G. Cazuguel, M. D. Abràmoff, B. Cochener, and C. Roux. Weakly supervised classification of medical images. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, pages 110–113. IEEE, 2012. [101](#)
- [83] S. Ravishankar, A. Jain, and A. Mittal. Automated feature extraction for early detection of diabetic retinopathy in fundus images. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 210–217. IEEE, 2009. [17](#), [18](#), [63](#)
- [84] P. Salembier. *Multiscale image analysis and modeling using rank order based filters - application to defect detection*. 1991. [20](#)
- [85] P. Salembier, A. Oliveras, and L. Garrido. Antiextensive connected operators for image and sequence processing. *Image Processing, IEEE Transactions on*, 7(4):555–570, 1998. [25](#), [90](#)
- [86] C. Sánchez, M. García, A. Mayo, M. López, and R. Hornero. Retinal image analysis based on mixture models to detect hard exudates. *Medical Image Analysis*, 13(4):650–658, 2009. [37](#)

BIBLIOGRAPHY

- [87] C. Sánchez, M. Niemeijer, M. Suttorp Schulten, M. Abràmoff, and B. van Ginneken. Improving hard exudate detection in retinal images through a combination of local and contextual information. In *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*, pages 5–8. IEEE, 2010. [37](#)
- [88] C. Sánchez, M. Niemeijer, I. Išgum, A. Dumitrescu, M. Suttorp-Schulten, M. Abràmoff, and B. van Ginneken. Contextual computer-aided detection: Improving bright lesion detection in retinal images and coronary calcification identification in ct scans. *Medical Image Analysis*, 16(1):50–62, 2012. [30](#), [31](#), [37](#), [48](#), [50](#), [53](#), [54](#)
- [89] C. I. Sánchez, R. Hornero, A. Mayo, and M. García. Mixture model-based clustering and logistic regression for automatic detection of microaneurysms in retinal images. In *SPIE medical imaging*, pages 72601M–72601M. International Society for Optics and Photonics, 2009. [63](#)
- [90] G. S. Scotland, P. McNamee, S. Philip, A. D. Fleming, K. A. Goatman, G. J. Prescott, S. Fonseca, P. F. Sharp, and J. A. Olson. Cost-effectiveness of implementing automated grading within the national screening programme for diabetic retinopathy in scotland. *British Journal of Ophthalmology*, 91(11):1518–1523, Jan. 2007. ISSN , 1468-2079. doi: 10.1136/bjo.2007.120972. [10](#)
- [91] A. Serna, B. Marcotegui, E. Decencière, T. Baldeweck, A.-M. Pena, and S. Brizion. Segmentation of elongated objects using attribute profiles and area stability. *Pattern Recognition*, 2014. Manuscript submitted for publication. [24](#)
- [92] J. Serra. *Image analysis and mathematical morphology*. London.: Academic Press. Mathematics, 1982. [21](#)
- [93] J. Serra, editor. *Image Analysis and Mathematical Morphology - Volume II : Theoretical Advances*. Academic Press, London, 1988. [119](#)
- [94] C. Sinthanayothin, J. Boyce, T. Williamson, H. Cook, E. Mensah, S. Lal, and D. Usher. Automated detection of diabetic retinopathy on digital fundus images. *Diabetic Medicine*, 19(2):105–112, 2002. [17](#), [37](#), [63](#)
- [95] P. Soille. *Morphological image analysis: principles and applications*. Springer-Verlag New York, Inc., 2003. [21](#)
- [96] A. Sopharak, B. Uyyanonvara, S. Barman, and T. Williamson. Automatic detection of diabetic retinopathy exudates from non-dilated retinal images using mathematical morphology methods. *Computerized Medical Imaging and Graphics*, 32(8):720–727, 2008. [37](#)
- [97] A. Sopharak, B. Uyyanonvara, and S. Barman. Automatic exudate detection from non-dilated diabetic retinopathy retinal images using fuzzy c-means clustering. *Sensors*, 9(3):2148–2161, 2009. [37](#)
- [98] T. Spencer, J. A. Olson, K. C. McHardy, P. F. Sharp, and J. V. Forrester. An image-processing strategy for the segmentation and quantification of microaneurysms in fluorescein angiograms of the ocular fundus. *Computers and biomedical research*, 29(4):284–302, 1996. [62](#), [63](#), [117](#)
- [99] J. Staal, M. D. Abràmoff, M. Niemeijer, M. A. Viergever, and B. van Ginneken. Ridge-based vessel segmentation in color images of the retina. *Medical Imaging, IEEE Transactions on*, 23(4):501–509, 2004. [118](#)

- [100] S. R. Sternberg. Grayscale morphology. *Computer Vision, Graphics, and Image Processing*, 35(3):333–355, 1986. [21](#), [119](#)
- [101] L. Streeter and M. J. Cree. Microaneurysm detection in colour fundus images. *Image Vision Comput. New Zealand*, pages 280–284, 2003. [63](#)
- [102] L. Tang, M. Niemeijer, J. M. Reinhardt, M. K. Garvin, and M. D. Abramoff. Splat feature classification with application to retinal hemorrhage detection in fundus images. 2013. [30](#), [84](#)
- [103] K. Tobin, E. Chaum, V. Govindasamy, and T. Karnowski. Detection of anatomic structures in human retinal imagery. *Medical Imaging, IEEE Transactions on*, 26(12):1729–1739, 2007. [40](#), [129](#), [133](#)
- [104] Y. A. Tolia and S. M. Panas. A fuzzy vessel tracking algorithm for retinal images based on fuzzy clustering. *Medical Imaging, IEEE Transactions on*, 17(2):263–273, 1998. [118](#)
- [105] D. Usher, M. Dumskyj, M. Himaga, T. Williamson, S. Nussey, and J. Boyce. Automated detection of diabetic retinopathy in digital retinal images: a tool for diabetic retinopathy screening. *Diabetic Medicine*, 21(1):84–90, 2004. [37](#), [63](#)
- [106] L. M. Vincent. Efficient computation of various types of skeletons. In *Medical Imaging V: Image Processing*, pages 297–311. International Society for Optics and Photonics, 1991. [132](#)
- [107] T. Walter and J.-C. Klein. Segmentation of color fundus images of the human retina: Detection of the optic disc and the vascular tree using morphological techniques. In *Medical Data Analysis*, pages 282–287. Springer, 2001. [117](#), [119](#), [129](#)
- [108] T. Walter, J. Klein, P. Massin, and A. Erginay. A contribution of image processing to the diagnosis of diabetic retinopathy-detection of exudates in color fundus images of the human retina. *Medical Imaging, IEEE Transactions on*, 21(10):1236–1243, 2002. [37](#), [43](#)
- [109] T. Walter, P. Massin, A. Erginay, R. Ordonez, C. Jeulin, and J.-C. Klein. Automatic detection of microaneurysms in color fundus images. *Medical Image Analysis*, 11(6):555, 2007. [17](#), [18](#), [19](#), [23](#), [63](#)
- [110] C. Wolf and J.-M. Jolion. Object count/area graphs for the evaluation of object detection and segmentation algorithms. *International Journal of Document Analysis and Recognition (IJ DAR)*, 8(4):280–296, Sept. 2006. ISSN 1433-2833, 1433-2825. doi: 10.1007/s10032-006-0014-0. [53](#)
- [111] D. Wu, M. Zhang, J.-C. Liu, and W. Bauman. On the adaptive detection of blood vessels in retinal images. *Biomedical Engineering, IEEE Transactions on*, 53(2):341–343, 2006. [17](#)
- [112] G. Yang, L. Gagnon, S. Wang, and M. Boucher. Algorithm for detecting micro-aneurysms in low-resolution color retinal images. 2001. [63](#)
- [113] A. A. Youssif, A. Z. Ghalwash, and A. S. Ghoneim. Comparative study of contrast enhancement and illumination equalization methods for retinal vasculature segmentation. In *Proc. Cairo Int. Biomed. Eng. Conf*, pages 21–24, 2006. [17](#)
- [114] B. Zhang, X. Wu, J. You, Q. Li, and F. Karray. Detection of microaneurysms using multi-scale correlation coefficients. *Pattern Recognition*, 43(6):2237–2248, 2010. [63](#)

BIBLIOGRAPHY

- [115] X. Zhang and O. Chutatape. Top-down and bottom-up strategies in lesion detection of background diabetic retinopathy. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 422–428. IEEE, 2005. [84](#)
- [116] X. Zhang and G. Fan. Retinal spot lesion detection using adaptive multiscale morphological processing. In *Advances in Visual Computing*, pages 490–501. Springer, 2006. [84](#), [98](#)
- [117] X. Zhang, G. Thibault, and E. Decenci ere. application of the morphological ultimate opening to the detection of microaneurysms on eye fundus images from clinical databases. *ICS*, 2011. [63](#), [108](#)
- [118] X. Zhang, G. Thibault, and E. Decenci ere. Automatic detection of exudates in color retinal images. 2012. [108](#)
- [119] X. Zhang, G. Thibault, and E. Decenci ere. Proc ed e de normalisation d echelle d’images ophtalmologiques (patent, filing number: 12 53929). "Patent", 2012. [29](#), [108](#)
- [120] X. Zhang, G. Thibault, E. Decenci ere, G. Qu ellec, G. Gazuguel, A. Erginay, P. Massin, and A. Chabouis. Spatial normalization of eye fundus images. In *International Symposium on Biomedical Imaging - ISBI*. IEEE, 2012. [108](#)
- [121] X. Zhang, G. Thibault, and E. Decenci ere. Exudate detection in color retinal images for mass screening of diabetic retinopathy. *Medical Image Analysis*, 2014. [108](#)
- [122] K. Zuiderveld. Contrast limited adaptive histogram equalization. In *Graphics gems IV*, pages 474–485. Academic Press Professional, Inc., 1994. [17](#), [18](#), [19](#), [23](#)

Image processing methods for computer-aided screening of diabetic retinopathy

Abstract : Diabetic retinopathy is the main cause of blindness among the middle-aged population. An early detection and adapted treatment considerably reduce the risk of sight loss. Medical authorities recommend an annual examination to diabetic patients. Several diabetic retinopathy screening programs have been deployed to enforce this recommendation. The aim of the TeleOphta project was to automatically detect normal examinations in a diabetic screening system, in order to reduce the burden on readers, and therefore serve more patients. This thesis proposes several methods to extract information linked to diabetic retinopathy lesions from color eye fundus images.

The detection of exudates, microaneurysms and hemorrhages is discussed in detail. One of the main challenges of this work is to deal with clinical images with high heterogeneity. New pre-processing methods which are capable to detect reflections and artifacts in the images, are proposed. Novel candidate segmentation methods, and new textural and contextual features for lesion characterization, are proposed. A random forest algorithm is used to detect lesions among the candidates. The proposed methods make extensive use of new residue analysis methods.

Moreover, three new publicly available retinal image databases, e-ophta EX, MA and HM, respectively designed to develop and evaluate exudate, microaneurysms and hemorrhages detections methods, are proposed in this work. Manual annotations of the lesions are given in detail.

The proposed methods have been integrated within the TeleOphta system, which is evaluated on two large databases. Each patient record is classified into two categories : "To be referred" or "Not to be referred". The classification is based on the results of the presented methods and on image signatures provided by other partners, as well as on medical information. The evaluation shows that the TeleOphta system can make about 2 times more patients benefit from the diagnosis service, based on the existing telemedicine network.

Keywords : Image processing, mathematical morphology, pattern recognition, diabetic retinopathy screening.

Méthodes de traitement d'images pour le dépistage de la rétinopathie diabétique assisté par ordinateur

Résumé: La rétinopathie diabétique est la cause principale de cécité dans la population en âge de travailler. Une détection précoce et un traitement adapté permettent de réduire considérablement le risque de perte de vue. Les autorités médicales recommandent un examen annuel pour les patients diabétiques. Plusieurs programmes de dépistage de la rétinopathie diabétique ont été déployés pour appliquer cette recommandation. L'objectif du projet TeleOphta était de détecter automatiquement des examens normaux dans un système de dépistage du diabète, afin de réduire le fardeau des lecteurs, et donc servir plus de patients. Cette thèse propose plusieurs méthodes pour extraire des informations liées à des lésions provoquées par la rétinopathie diabétique dans des images en couleurs du fond d'œil.

La détection des exsudats, microanévrismes et hémorragies est discutée en détail. L'un des principaux défis de ce travail est de traiter des images cliniques avec haute hétérogénéité. Des nouvelles méthodes de pré-traitement, qui effectuent non seulement des tâches de normalisation et de débruitage, mais aussi de détection de réflexions et d'artefacts optiques, sont proposées. Des méthodes de segmentation des candidats basées sur la morphologie mathématique, et de nouveaux descripteurs de texture et de contexte sont proposées pour la caractérisation des lésions. Un algorithme de forêts aléatoires est utilisé pour choisir les lésions parmi les candidats. Les méthodes proposées utilisent largement des nouvelles méthodes d'analyse des résidus.

En outre, trois nouvelles bases de données publiques d'images de la rétine, e-ophta EX, MA et HM, respectivement conçues pour développer et évaluer les méthodes de détection d'exsudats, de microanévrismes et d'hémorragies, sont proposées dans ce travail. Des annotations manuelles détaillées des lésions sont fournies.

Les méthodes proposées ont été intégrées dans le système TeleOphta, qui est évaluée sur deux grandes bases de données. Chaque dossier du patient est classé en deux catégories: "Pour avis" ou "Normal". La classification est basée sur les résultats des méthodes présentées et sur les signatures d'image fournies par d'autres partenaires, ainsi que sur l'information médicale du patient. L'évaluation montre que le système TeleOphta permet de traiter deux fois plus de patients dans un réseau de dépistage, à moyens constants.

Mots clés: Traitement d'images, morphologie mathématique, reconnaissance de formes, dépistage de la rétinopathie diabétique.

