Robust Eye Biometry and Genetics

Stine Harder & Susanne Rytter Christoffersen

Kongens Lyngby 2012
IMM-M.Sc.-2012-0008

Technical University of Denmark
Informatics and Mathematical Modelling
Building 321, DK-2800 Kongens Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk IMM-M.Sc.-2012-0008

Summary (English)

The goal of this thesis is to make an objective evaluation of the human iris colour and structure, in order to extract quantitative measures for DNA correlation.

In order to fulfil this goal the iris region was extracted from high resolution eye images provided by the Department of Forensic Medicine, Section of Forensic Genetics. Inspired by a method proposed by John G. Daugman we have succeeded in extracting the iris region. Furthermore the eyelid boundaries were located using a pixel classification with a Markov Random Field approach. We were able to locate the pupil boundary and iris boundary correctly in 75.0% and 96.4% of the cases and the upper and lower eyelid boundaries in 76.6% and 73.0% of the cases, respectively. The total percentage of correctly extracted irises is 64.3% using a Dice Coefficient of 0.92 for comparison with human annotations.

The evaluation of human eye colour was performed by generating a colour classifier, an explanatory ratio, and by performing a colour based image clustering. The colour classifier was trained on subjective evaluations of human eye colour, and the percentage of correctly classified images was 65.4%. We have calculated a ratio explaining the amount of blue and brown coloured regions in the iris, which has shown to be a promising continuous measure for iris colour since it clearly distinguishes between blue and brown eye colour. The colour based image clustering is a data driven approach, where the division into colour groups are very robust compared to a subjective evaluation. The image groups generated follow the general perception of eye colour being divided into groups ranging from blue, through intermediate, to brown. This is in agreement with the knowledge about DNA.

A minor study of human perception of eye structures was conducted, resulting in Fleiss' Kappa values between 0.23 to 0.56 for different structures.

A structure based image clustering was performed in order to generate a data driven evaluation of iris structures. The result is a division into three groups, one consisting of very smooth eye images, one consisting of eye images with a high amount of structure and one consisting of images with a large reflection area. This division is very promising regarding a DNA correlation. However some images were incorrectly clustered, which makes us question the result.

Summary (Danish)

Målet med denne afhandling er at udvikle en metode til objektiv evaluering af det menneskelige øjes farve og struktur, samt at udtrække kvantitative mål til brug ved korrelation med DNA.

For at opnå dette mål blev iris informationen udtrukket fra en række højt opløselige øjenbilleder udleveret af Retsmedicinsk Institut, Retsgenetisk Afdeling. Vi har, inspireret af en metode udviklet af John G. Daugman, succesfuldt udtrukket irisregionen. Derudover har vi været i stand til at lokalisere øjenlågsgrænserne på baggrund af en pixelklassificering, løst ved brug af Markov Random Fields. Vi har succesfuldt lokaliseret pupillen, iris, øvre og nedre øjenlåg i henholdsvis $75,0\%$, $96,4\%$, $76,6\%$ og $73,0\%$ af tilfældene. Den totale succesrate for hele irisudtrækningen er $64,3\%$ ved brug af en Dice koefficient på 0,92 som sammenligning med menneskelige annoteringer.

Evalueringen af øjenfarve blev foretaget ved at generere en farveklassificering, en forklarende ratio og ved at udvikle en farvebaseret billedgruppering. Farveklassificeringen blev trænet på baggrund af subjektive vurderinger af øjenfarve. Procentdelen af korrekt klassificerede billeder var 65.4%. Den udregnede ratio mellem blå og brune regioner i iris har vist sig at være et lovende mål for irisfarve, idet en klar skelnen mellem blå og brun øjenfarve er tilstede.

Den farvebaserede billedgruppering er en datadreven model, hvor opdelingen i farvegrupper er mere robust sammenlignet med den trænede farveklassificering. De resulterende farvegrupper følger den generelle opfattelse af øjenfarve, inddelt i grupper gående fra blå, over mellemliggende til brun.

Et mindre studie af den menneskelige opfattelse af øjenstrukturer blev udført med resultat i Fleiss' Kappaværdier på mellem 0,23 og 0,56 for forskellige strukturer.

En strukturbaseret gruppering blev udført for at generere en datadreven evaluering af iris strukturer. Resultatet er en opdeling af strukturer i tre grupper, en bestående af meget glatte øjenbilleder, en anden bestående af øjenbilleder med stor mængde struktur og en tredje bestående af øjenbilleder med store reflektionsområder. Denne opdeling er lovende i forbindelse med DNA korrelation. Dog skal det bemærkes, at nogle billeder bliver ukorrekt grupperet, hvilket får os til at tvivle på resultaterne.

Preface

This thesis was prepared in collaboration with the Department of Informatics and Mathematical Modelling at the Technical University of Denmark and the Department of Forensic Medicine, Section of Forensic Genetics. It fulfils the requirements for acquiring an M.Sc. degree in engineering for Medicine and Technology.

The thesis deals with extracting colour and structural information from the iris using high resolution eye images. The thesis is made on the basis of eye images received from the Department of Forensic Medicine, Section of Forensic Genetics. The main focus is to provide the Department of Forensic Medicine, Section of Forensic Genetics with relevant parametrisations of colour and structure of the iris.

The thesis consists of this assignment, implementation code, a graphical user interface (GUI) platform for evaluating iris colour and annotating structures and a poster presentation that was presented at the annual Medico Bazar arranged by Medico Innovation [1]. The code from the implementations, along with the GUI and the poster, is available on the CD-ROM provided with this thesis. The contents of the CD-ROM are seen in Appendix E.

Lyngby, 29-February-2012

Stine Harder Susanne R

Stine Harder & Susanne Rytter Christoffersen

[1] http://www.medico-innovation.dk/

Acknowledgements

We would like to thank The Department of Forensic Medicine, Section of Forensic Genetics for providing us with eye image data and for providing us with office space. A special thanks to Professor, DMSc Niels Morling, cand. scient Peter Johansen, cand. scient Jeppe Dyrberg Andersen and Senior Advisor, Ph.D. Claus Børsting for taking their time to meet with us and for helping us with the genetic aspects of the thesis. Furthermore, thanks to Anja Ladegård Jørgensen, Maibritt Sigvardt, Trine Leerhøj Hansen, Nadia Jochumsen, Peter Johansen and Jeppe Dyrberg Andersen for annotating and evaluating our validation images.

We would like to thank our supervisor Associate Professor Rasmus Reinhold Paulsen for his enthusiasm and many great recommendations regarding our project course. Thanks to the staff at Department for Mathematical Modelling at the Technical University of Denmark for a very warm welcoming at the department and at the Christmas workshop, and for asking critical questions.

A special thanks to Assistant Professor Anders Lindbjerg Dahl for providing us with code regarding the DAISY descriptor and for his great involvement in our thesis. Furthermore, thanks to Associate Professor Henrik Aanaes for providing us with code regarding Markov Random Fields during the course 02503, Advanced Image Analysis, spring 2011.
Thanks to Associate Professor Karl Sjöstrand and PhD student Rasmus Ramsbøl Jensen for supervision on Principal Component Analysis and Markov Random Fields, respectively.

Finally thanks to those who have read and commented on our thesis.

Contents

Chapter 1

Introduction

1.1 Human Genetic Evolution and Physical Traits

If looking at humans across the world it is easily appreciated that several differences in appearance exist. Take for example, the differences between ethnic populations such as Europeans, Asians, Native Americans and Africans. Historically seen, the ongoing evolution of the modern human originates 200.000 years back to the geographical region of Africa. In this region of the world most genetic variations are found, and research has shown that the geographical expansion of humans has its origin on this continent [4]. The genetic variations among populations is a result of adaptation to a constantly changing environment and selective events occurring over the last 200.000 years. These events have resulted in demographic changes and the geographical expansion of humans into other parts of the world, creating different ethnic populations around the world [4].

Each ethnic population has different traits contained in the population, but in general all humans differ by having unique traits. Scientists around the world have for many years made a successful effort to define the human evolution on a molecular level. The unique traits of every human being are found to be mainly specified through genetic variations that are inherited through generations [1]. This knowledge has been used in several research areas, for example to identify

inherited diseases, identify ancestry, predict health issues, and so forth.

The visual uniqueness of humans is expressed through the physical traits, such as hair colour, skin colour, eye colour, height, weight and several other traits. Such visible traits are called external visible characteristics (EVC) [35, 17].

1.2 Forensic Genetics

In the field of forensic genetics, the main goal is to identify a person from a sample of genetic material. The main usage of identification procedures using deoxyribonucleic acid (DNA) material is constituted in paternity cases, immigration cases and crime cases. The interesting field regarding this thesis is the use of DNA for crime cases. When a crime is committed, the police authority collects different kinds of effects from the crime scene. If present, part of the effects might encompass DNA material. This DNA sample is analysed and compared to the criminal DNA register. If no match is found, there is a need to find a suspect in order to make a comparison between the suspect DNA and the DNA sample from the crime scene.

This thesis is the first step towards generating an objective tool for the authority to use in cases where there is no suspect to compare to the DNA sample. The idea is to be able to predict some EVCs of the suspect based on the DNA from the crime scene. It would be a great tool for the police to be able to predict for example entire facial characteristics of the suspect. This would help the police by either restricting a large pool of suspects or by confirming human eyewitness statements.

1.2.1 Estimation of EVCs

The road towards predicting appearance from a DNA sample is long, since finding the link between an EVC and genetics is not an easy task. One of the main difficulties when estimating EVCs from genetic material is that the traits are very complex on a genetic level, due to the fact that most EVCs are inherited as a polygenic trait. Thus several genes constitute the phenotype [12]. In general the human facial characteristics are of great interest when identifying people. One very interesting EVC is the colour and structure of the human eye. Since the human perception of iris colour is a combination of the colour and the underlying structures, both colour and structures are valuable traits if identifiable in the genome.

Much of the current research on linking genetic material to iris colour is based on subjective classifications of the colour content, performed by qualified experts in the field of eye research [35, 23]. The subjective classifications are compared with genetic material and in this way the genes that constitute the iris colour are identified. This result in a classification that cannot be replicated, since people perceive colour differently. A main interest is therefore to create an objective method to classify iris colour and structure. An objective method can improve the identification of genetic material by omitting subjective classification that cannot be validated by third parties.

To the best of our knowledge, no previous research has focused on the subject of linking iris structures to genetic material, why the subject is a completely new area of research.

Looking at the human eye, the perceived EVC is the colour of the iris. The iris colour is the most visible characteristic of the eye, but when investigating the iris further it is seen that even more physical traits are present in the structure of the iris. In order to understand the colour perception of the iris a short introduction is presented.

1.2.1.1 The human iris

The outer anatomy of the eye are seen in Figure 1.1 which shows an example of an eye image. The arrows indicate specific structures of the outer anatomy of the eye, along with a reflection from the camera flashes. These reflections are present in all of our eye image data.

The general conception of iris colour is that the perceived colour is ranging from the lightest blue to the darkest brown, where the shades between blue and brown is a result of limited pigmentation content, iris thickness and cell density [32]. However, this conception is not widely used due to the highly subjective assessment of iris colour that makes it difficult to distinguish iris colour into a high number of categories. A scale of blue, green-hazel and brown is therefore often used [32]. Seen from a genetic point of view these groups have a specific construction of the DNA, also called the genotype.

The complex tissue structures of the iris have been investigated by many researchers and a series of different characteristic structures are defined. In general four widely used structures are: Fuchs Crypt's, Nevi dots, Wolfflin nodules and so-called contraction furrows, see Figure 1.2 for an illustration. A Fuchs Crypt is described as "recessed folds of the iris tissue radiating out from the pupil" [26], contraction furrows are described as "circular folds around the iris" [26],

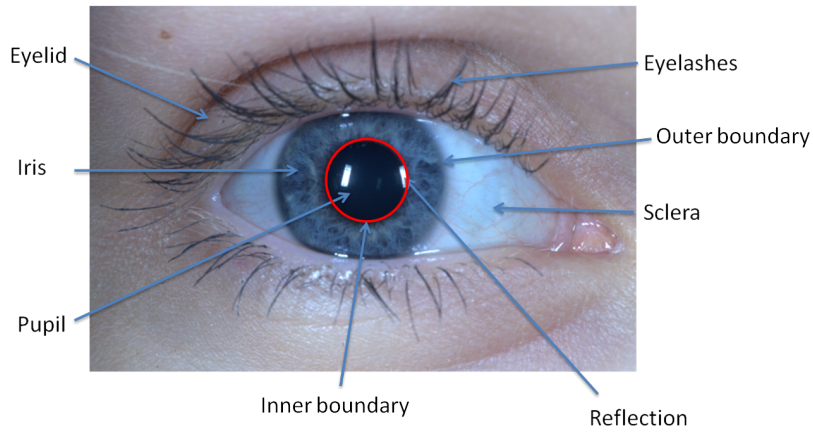

Figure 1.1: Illustration of the outer anatomy of the human eye

Wolfflin nodules are described as "lightly pigmented spots around the periphery of the iris" [26] and Nevi dots are generally defined as pigmentation placed on the surface of the iris structures [32]. These specific structures constituting the unique biometric makers of the iris are what makes it possible to think of the iris as a fingerprint.

1.2.2 Goal of the thesis

From the previous sections it is clear that there is a need for a thorough investigation of the iris colour and structures in order to predict colour and structure from a DNA sample. A collaboration between the Department for Mathematical Modelling at the Technical University of Denmark and the Department of Forensic Medicine, Section of Forensic Genetics makes it possible to start solving the task. Our thesis is the first step in the process, and the main result from the thesis should be an objective and quantitative evaluation of the colours and structures of the iris. Subsequently the Department of Forensic Medicine, Section of Forensic Genetics will make a correlation between the objective measurements extracted from the iris and the corresponding DNA sample.

For this thesis a number of eye images were given as will be described in Chapter 2. These images contain skin, sclera, eyelashes, eyelids and of course the iris and the pupil. The interesting part of the images is the iris area, and therefore the first part of this thesis will be concerned with extracting the iris region. The task of extracting the iris region can be partitioned into two major steps. The first step is to locate the iris boundaries, as explained in Section 3.1. The

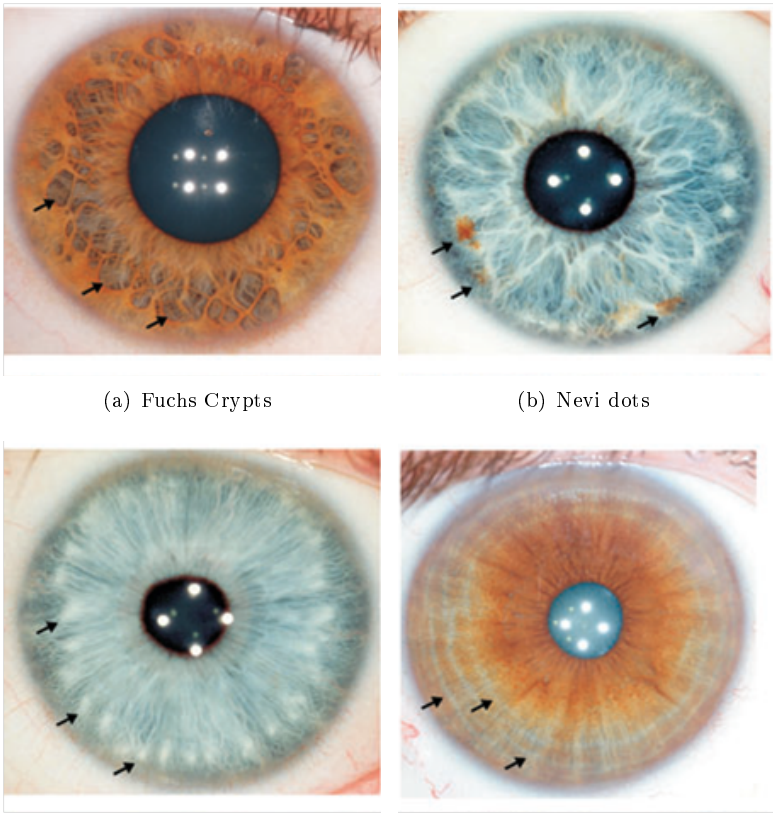

(a) Fuchs Crypts (b) Nevi dots

(c) Wolfflin nodules (d) Contraction furrows

Figure 1.2: a) Arrows indicating Fuchs Crypts b)Arrows indicating nevi dots c) Arrows indicating Wolfflin nodules d) Arrows indicating contraction furrows. All four images are taken from [32]

second step is to locate the eyelid boundaries, since the eyelids often occlude the iris area, see Section 3.3. However this can be a difficult task, since the eyelashes often occlude the eyelid boundaries and since a high variation in eye colour exist.

In order to make a quantization of the iris colour a colour classification method is proposed. The colour classifier should be able to classify an eye image into the same class every time, so that results can be replicated. The colour classifier is explained in Chapter 4. The objective colour classifier is a step in the right direction towards obtaining a reliable parametrization of the iris colour for correlated with genetics.

Another interesting quantization measure could be a ratio describing the amount of blue and brown in the eye images. Our belief is that a ratio between the brown and the blue areas in an eye would be an appropriate measure, see Chapter 5.

A goal of this thesis is also to investigate the various structures in the iris. This is done by examining the operator variance between humans for a number of predefined structures, see Chapter 6.

The approach used for investigating the structures in the iris is a data driven clustering approach, see Chapter 7. This approach has multiple goals. The first goal is to be able to locate groups of irises with similar colour or structure, which could be correlated with DNA samples by the Department of Forensic Medicine, Section of Forensic Genetics in order to investigate the genetics. The second goal is to investigate whether the amount of structures is dependent on the iris colour, which is of interest for the Department of Forensic Medicine, Section of Forensic Genetics.

Chapter 2

Data Acquisition and Preprocessing

The image data was acquired from the Department of Forensic Medicine, Section of Forensic Genetics, who in collaboration with Glostrup Hospital and Bispebjerg Hospital has taken a wide span of high resolution digital images of patient eyes. The images were taken with the following camera and camera settings; Canon Eos 5D Mark II, Canon Macro Lens 100mm 67Ø. 1:2.8 L IS USM, Canon Macro Twin Lite MT-24EX,ISO 800, Shutter 1/100, AV 18. The data were saved in raw format with a size of 5634×3753 pixels. An illustration of the camera is seen in Figure 2.1.

The images were taken at the hospitals, by employees from the Department of Forensic Medicine, Section of Forensic Genetics. They were acquired by a hand held camera as illustrated in Figure 2.2. For an optimal image of the eye it is important that the patient concentrates on keeping the eyes wide open, thereby minimizing occlusion of the iris from the eyelids and eyelashes. This is, however, a demanding task, and the patients were therefore instructed not to physically force their eyes open.

Preprocessing of the images were done in order to make the image size compatible with MATLAB. The images were converted to white balanced JPG images,

(a) Camera viewed from the front (b) Camera viewed from the side

Figure 2.1: The camera used for image acquisition

Figure 2.2: Photo of the camera in use. The image illustrates the camera position during image acquisition

using a Canon provided software, the Digital Photo Professional (DPP)[2]. A conversion of the images to reduce the size of the Portable Network Graphics (PNG) images, using the AnyPic Image Resizer Pro program[3] was performed. The reduced image size was 639×426 pixels. A total of 275 images were provided for the project. Since many of the tasks during the thesis needs both training images and test images, and because we needed a final group of images for validation, the 275 images were divided into three groups. The group for training contains 100 eye images, the group for testing contains 64 eye images and the group for validation contains the remaining 111 eye images. The three groups of eye images are from now on referred to as training data, test data and validation data, respectively.

[2] Available from http://www.canon.dk/For_Home/Product_Finder/Cameras/Digital_SLR/eos30d/software/

[3] Available from http://www.batchimageconverter.com/image-resizer-pro/

Chapter 3

Iris Extraction

In order to analyse the iris, the iris region has to be located and extracted from the eye images. When looking at an eye images, two iris boundaries can be defined, one for the boundary of the pupil and one for the boundary of the iris, see Figure 1.1. These boundaries will throughout this thesis be denoted as the inner boundary and outer boundary of the iris.

Previously, other researchers have extracted the iris region from an image, mostly in order to perform biometric iris matching. The leading researcher in the area is John G. Daugman, who developed a biometric person identification system based on iris analysis. Our approach for extracting the iris is in some aspects very similar to the method of John G. Daugman. The basic idea for iris extraction is to fit a circle to the gradients of the inner and outer iris boundaries [5, 7].

However, for our eye images fitting a circle is not enough, since the iris is often occluded by the eyelids. Looking at Figure 1.1 it is seen that the outer boundary is occluded by the upper and lower eyelids, which is often the case. The occlusion from the eyelids adds another task to the iris extraction procedure, namely the detection of the eyelid boundaries. Furthermore, our eye images contain reflections from the camera flashes, see Figure 1.1, which should also be accounted for.

This chapter on iris extraction consists of a short explanation of the method proposed by John G. Daugman, upon which our implementation of the method will be explained. Our implementation consists of several steps: the inner and outer boundary detection, a radial transformation of the iris region, a description of how we account for the reflections from the camera flashes, and a description of how we detect the eyelid boundaries.

3.1 Inner and Outer Boundary Detection

Our method for detecting the inner and outer iris boundaries is inspired from the method proposed by John G. Daugman. In order to ease the iris extraction procedure, the eye images are converted into gray scale images.

3.1.1 Iris extraction - John G. Daugman

John G. Daugman has developed a very known and acknowledged method for iris extraction [5, 7]. He takes advantage of the fact that the inner and outer boundaries are circular and that there is a difference in pixel intensity between the pupil and the iris and between the iris and sclera regions. John G. Daugman's method is therefore to find the highest sum of gradients on a circle.

The method of John G. Daugman is based on:

$$\max(r, x_0, y_0) \left| G_\sigma(r) * \frac{\partial}{\partial r} \oint_{r,x_0,y_0} \frac{I(x,y)}{2\pi r} ds \right| \tag{3.1}$$

where $I(x, y)$ is an image, ds is a circular arc with radius r and center coordinates (x_0, y_0). The symbol $*$ denotes convolution and $G_\sigma(r)$ is a Gaussian smoothing function with scale parameter σ [5, 7].

Equation 3.1 basically states that the inner and outer boundaries are found by searching the eye image for the highest sum of gradients on a circle. The search is performed by changing the center coordinate and radius. The sum of gradients is calculated on a Gaussian smoothed image, since an eye image contains many small structures that could be detected as circular edges otherwise.

According to John G. Daugman both inner and outer boundary of the iris should be located using Equation 3.1. When the iris boundaries have been extracted,

John G. Daugman suggests a similar approach in detecting eyelid boundaries. The eyelid detection is based on curvilinear edges, and the proposal by John G. Daugman is to sum along an arc instead of a circle. An example of his result is seen in Figure 3.1, where it is seen that the iris region has been isolated from other image regions.

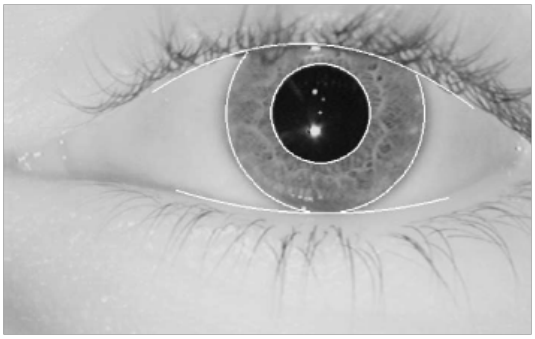

Figure 3.1: Result from iris extraction by John G. Daugman. Taken from [5]

John G. Daugman has developed a biometric iris measure for person identification. The result of his work therefore focuses on the quality of the iris matching and not on the quality of the detection of the iris boundaries. The accuracy of the boundary detection has therefore not been found reported.

3.1.2 Inner and outer boundary detector

In this section our method for detecting the inner and outer iris boundaries will be explained. Our method for iris extraction is based on the method proposed by John G. Daugman [5] [7], although some differences are present. The method proposed by John G. Daugman is highly optimized according to performance time, which complicates the implementation. Our approach is more straight forward, which is possible since the time aspect is not critical for our implementation.

Our implementation of the method proposed by John G. Daugman follows the basic idea very closely, while avoiding the more complicated optimizations performed by John G. Daugman [7]. As explained in Section 3.1.1 the basic idea is to find the two maximum sums of gradients on a circle, corresponding to the inner and outer iris boundaries, respectively. In order to avoid noise components in the image, such as small structures in the iris, the summation of gradients is performed on a Gaussian smoothed image.

The first step of the implementation is to make a Gaussian smoothing of the eye image. The probability density function of a Gaussian is expressed as:

$$g(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp^{-\frac{x^2-\mu^2}{2\sigma^2}} \tag{3.2}$$

where x is the data, μ is the mean and σ is the standard deviation [25].

Our implementation uses a Gaussian filter with a standard deviation $\sigma = 3$, see Figure 3.2. Multiple values for standard deviation have been tested with the algorithm, which has shown that the choice is critical for the implementation. In order to keep the performance of the algorithm high, the standard deviation should be in the range $\sigma \in [2.5, 3.5]$.

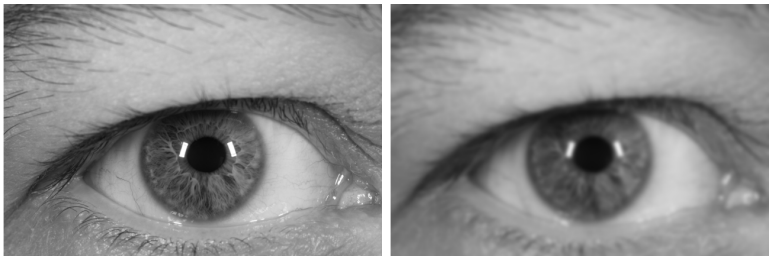

Figure 3.2: Left: Gray tone eye image. Right: Gaussian smoothed image, $\sigma = 3$

After performing Gaussian smoothing of the image, the sum of gradients on a circle should be calculated. This task can be divided into three steps: first find the sample positions on the circle, then calculate the gradient value for each position and finally make a summation of the values.

The positions (x, y) on a circle for a given center coordinate (x_0, y_0) and radius r is defined as:

$$x = r\cos(\Delta\theta p + 90) + x_0 \tag{3.3}$$
$$y = r\sin(\Delta\theta p + 90) + y_0 \tag{3.4}$$

where $\Delta\theta$ is the angular step size and $p \in [0, \frac{360}{\Delta\theta}]$. If a position on the circle is located between pixels, a bilinear interpolation is performed.

The gradient value should be calculated for every sample position on the circle. Since the gradients are calculated in different directions dependent on the position on the circle, an analytical differentiation cannot be made and instead a finite differentiation is used. The gradient value $\frac{dg}{dr}$ for a given position (x, y) can be calculated as:

$$\frac{dg}{dr} = \frac{f(r+1) - f(r)}{1} \tag{3.5}$$

where $f(r)$ is the pixel value for a given radius r and $f(r+1)$ is the pixel value one radial step away from r.

Finally a summation of gradient values, G, should be calculated. The summation of gradient values is given by:

$$G = \sum_{p} g(x, y) \tag{3.6}$$

where (x, y) are the sample positions on a circle given by Equation 3.3 and 3.4.

The gradient sum for a circle can now be calculated and the final step in the process of locating the boundaries is to find the maximum gradient sum by changing the center coordinate (x_0, y_0) and radius r:

$$\max\left(G(r, x_0, y_0)\right) \tag{3.7}$$

3.1.2.1 Optimization of processing time

Locating the circle with the highest gradient sum (e.g. the pupil boundary) can be a computational and time consuming task if gradient sums should be calculated for all image possible center coordinates and radii. Therefore it would be valuable to find a method that optimizes the processing time but still keeps the performance high. In this section two methods are explained, one for optimizing the search for center coordinate and one for optimizing the search for radius. Both methods are implemented into the algorithm.

Center coordinate search The task of locating the center coordinate is a task of locating a global maximum. The Steepest Descent method [10] would be an obvious choice, since it is a fast method to locate a local minimum or maximum. However, since the search function is not smooth, see Appendix A, the Steepest Descent method is not appropriate. Another method is therefore sought.

The method used in our algorithm reduces the processing time by calculating the gradient sums for a reduced amount of pixels. The idea is to start by calculating the gradient sums for a small number of pixels and then narrow the region of interest to an appropriate area around the pixel with the highest gradient sum. The procedure should then iterate until a small region of interest is located, and the gradient sum would then be calculated for all remaining pixels. In this way one pixel accuracy is obtained and computational effort and time is reduced.

For our implementation the horizontal and vertical step sizes are chosen as $\frac{\text{number of columns}}{10}$ and $\frac{\text{number of rows}}{10}$. After each iteration the search region is decreased to a fourth of the previous size of the search region. This was chosen in order to lower the processing time while maintaining the detection rate. Since the step size between the pixels is dependent on the size of the region of interest, the step size is also decreased. After each iteration, the pixel with the maximum gradient sum is found, and the search region is confined to a region around this coordinate. For the last iteration the gradient sum is calculated for all pixels in the small region of interest. The method is illustrated in Figure 3.3 a). The result of locating the outer iris boundary is seen in Figure 3.3 b). Obviously the outer boundary has been located, and the processing time has been decreased drastically.

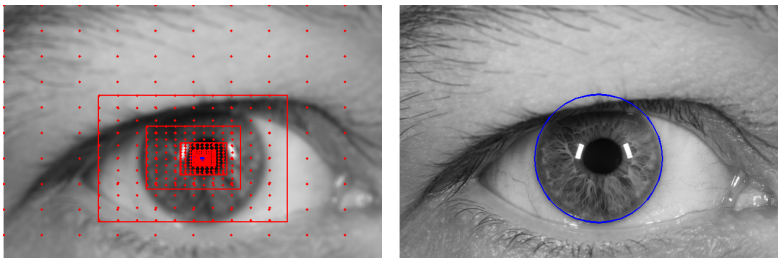

(a) Illustration of method (b) Result of method

Figure 3.3: a) Illustration of the center coordinate search for the outer boundary. The red dots indicate the sample pixels, the red squares indicate the region of interest and the blue dot indicate the located center coordinate. b) Result from locating the outer iris boundary

Radius search The approach used for optimizing the processing time with respect to the radius search is very alike the method used for the center coordinate search. The method is to calculate the gradient sum for a small number of radii, illustrated as the red circles in 3.4 a), and then narrow down the search area, illustrated as the magenta circles in 3.4 b). The circular gradient sum is then calculated for all radii in this restricted search area.

The radial step size is 10 pixels in the first run, and the new search area is chosen so that there are 10 pixels on both sides of the circle with maximum gradient sum. Figure 3.4 a) and b) illustrate the method, and c) shows the result when locating the inner boundary radius for the correctly given center coordinate.

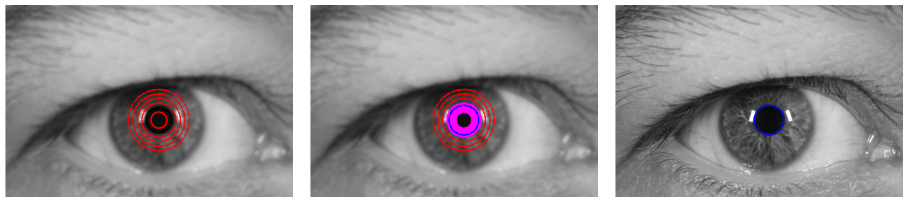

(a) First run (b) Second run (c) Result

Figure 3.4: Illustration for radius search for inner boundary. Red circles indicate the sample radii, black circle indicates radius with highest gradient sum, magenta circles indicate the sample radii for the second run, blue circle is the final result

3.1.2.2 Restrictions for inner and outer boundary

Eye images contain many gradients beside the gradients for the inner and outer iris boundaries. Very dark eyelashes is an example of a structure that creates large gradient values. These large gradients make the task of extracting the iris much more difficult. In order to make the implementation more robust, appropriate constraints are made on the size of the inner and outer boundaries.

According to [27], the diameter of the pupil varies from 1.5 mm to approximately 9 mm in a range from direct light to total darkness, respectively. We chose a minimum pupil diameter of 1.44 mm and a maximum pupil diameter of 8 mm. The reason for choosing such a large span of the pupil size is that our dataset contains eye images with a pupil diameter in the entire range. Literature on the subject of iris size is almost non-existing, but according to [20], the normal iris diameter is 12 mm. We chose a minimum iris diameter of 11.44 mm and an upper iris diameter of 15 mm. The choice of limits for the inner and outer

boundaries is very critical for the implementation. These limits were chosen since they gave the highest percentage of correctly detected boundaries for our eye images.

Another restriction is that the outer boundary is located first, and based on this result the search area for the inner boundary is restricted to be contained inside the outer boundary. This restriction is implemented since the pupil cannot physically be located outside the iris region.

3.1.2.3 Search cones

By inspiration from John G. Daugman [5, 7], two 90 degree cones are used when searching for the outer iris boundary. This is done because the outer iris boundary is often occluded by the eyelids. The 90 degree cones are illustrated in Figure 3.5 along with the full circle used for the inner boundary search.

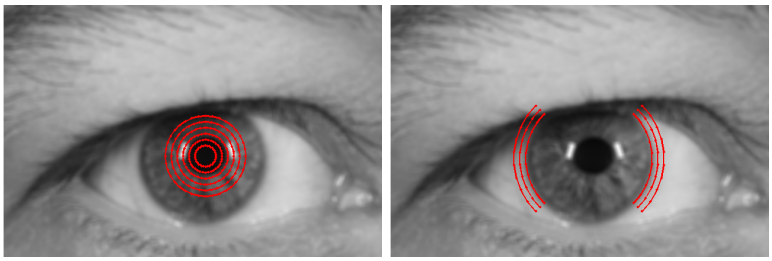

Figure 3.5: Left: Illustration of a full 360 degree cone. Right: Illustration of left and right 90 degree cones

3.1.2.4 Reflections from camera flashes

The eye image data received from the Department of Forensic Medicine, Section of Forensic Genetics, see Chapter 2, all contain reflections from the camera flashes, see Figure 1.1. The reflections are positioned either in the pupil region, the iris region or on the boundary between the pupil and iris.

Since the reflections contain very high image values compared to the surrounding pixels, a high gradient value is present at the reflection boundaries. The high gradients might lead to false detections of the inner and outer iris boundaries, and therefore the reflections must be accounted for. An example of incorrectly detected iris boundaries, due to reflections, is seen in Figure 3.6 a). Figure 3.6 b) shows the correctly detected iris boundaries.

Figure 3.6: Left: Result obtained without a mask to account for the reflections. Right: Result obtained with a mask accounting for the reflections

The approach used to account for the reflections is to generate a mask containing the image regions containing the reflections. By applying a simple threshold to the eye image, a mask is produced where all pixels with a value above the threshold are set to zero, and remaining pixels are one. The threshold was set to 240, which is very specific for our eye images. In order to assure that the entire reflection area found is represented in the mask, the located area is dilated.

The threshold and dilation procedure produces a mask with many areas defined as reflection areas, as seen in Figure 3.7 b). The original gray tone image is seen in a). In order to avoid e.g. one pixel sized regions, a limit is set on the size of the reflection area. The minimum size is set to 10 pixels, and the resulting mask is seen in Figure 3.7 c).

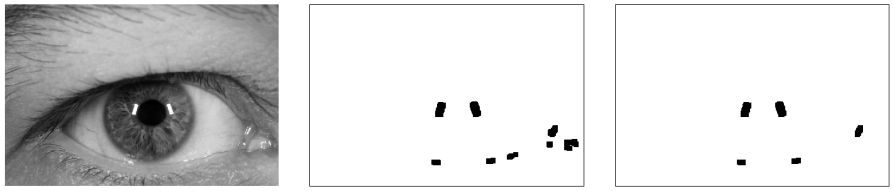

(a) Original gray tone image (b) Temporary mask (c) Final mask

Figure 3.7: a) Original gray tone image b) Temporary mask, created by simple threshold procedure c) Final mask, created by threshold procedure followed by a lower limit on the size of the regions

A binary reflection mask is generated for each eye image, which is used throughout the thesis in order to avoid using the information contained in the reflection areas.

3.2 Iris Mapping

When the inner and outer iris boundaries have been found, it is of interest to rearrange the information contained in the iris in a useful and logical manner. John G. Daugman [6] proposes the Rubber Sheet model, which is basically a radial transformation of the iris region.

For the radial transformation, points are sampled in the original eye image according to the location of the inner boundary. A specified number of points are sampled along a line from the inner boundary and outwards toward the outer boundary and mapped onto an iris map. The vertical axis on the iris map consists of a normalized radius, where 0 is at the inner boundary and 1 is at the outer boundary. The horizontal axis is the angular position, starting with 0 degrees at the bottom of the circle. The next line to sample is chosen by moving a specified angular step size around the circle. The new line is sampled in the same manner, and the procedure is continued until around the circle. The procedure is illustrated in Figure 3.8.

The remapping of the iris image $I(x, y)$ from raw coordinates (x, y) to dimensionless coordinates (r, θ) can according to [5] be represented as

$$I(x(r, \theta), y(r, \theta)) \rightarrow I(r, \theta) \tag{3.8}$$

The step size when sampling from the inner to the outer boundary, is changed for each angular step. The change in step size is performed because the distance from the inner to the outer boundary is not necessarily the same at all angular positions. This is due to the fact that the pupil and the iris center coordinates are not always located in the same position.

An example of a radial transformation of an eye image is seen in Figure 3.9.

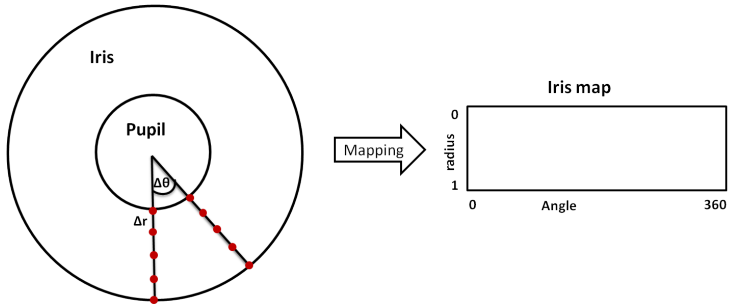

Figure 3.8: Illustration of the radial transformation. Left side shows the inner and outer iris boundaries, as well as the sampling steps, $\triangle$r, along the radial direction and the angular step $\triangle\theta$

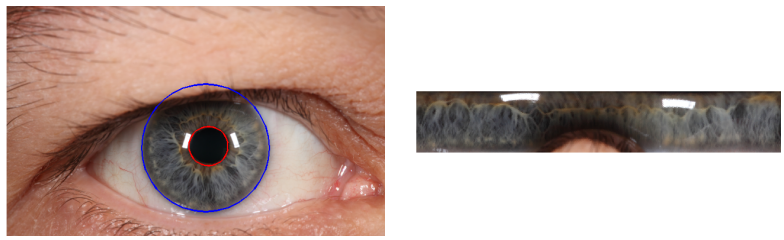

Figure 3.9: Left: Eye image with located inner and outer iris boundaries. Right: Iris map, which is the result from the radial transformation

3.3 Eyelid Detection

When looking through our database of eye images it is seen that most of the images contain an iris region that is occluded by one or both eyelids, see Figure 3.9 a) for an example. This occlusion will result in wrong information in the iris map resulting from the radial transformation, see Figure 3.9 b). This is a problem since the eyelid regions interfere with e.g. the mean eye colour and many other features. Locating the eyelid boundaries is therefore an important task.

The eyelid boundary detection has proven to be a difficult task. One of the main difficulties is that the eyelashes interfere with the boundaries between eyelid and inner eye area (sclera, iris and pupil). In order to detect the eyelid boundaries many different approaches have been tried, but most of them without promising results. However, we have succeeded in generating a method for eyelid detection.

3.3.1 Colour representations

The opportunity to change between colour representations is important, regarding both eyelid detection and other tasks throughout the thesis. Different colour representations can be useful for different tasks, and one colour representation might be appropriate for one task, but not for another. The RGB (red, green, blue) and HSV (hue, saturation, value) colour representations are often used in image processing, and the Commision Internationale de L'Eclairage $L^*a^*b^*$ (CIELAB) colour representation has been used for iris colour quantification in [9]. The basic understanding of the three colour representations is important, and the process of converting between the representations will be skipped, since the conversion of RGB to HSV and back can be performed by incorporated functions in MATLAB, and the RGB to CIELAB and back can be performed using the software by Mark Ruzon[4].

The eye images are captured using RGB values, hence every single pixel has a value for R, G and B. Figure 3.10 a) shows the RGB colour space, where all colours are represented as a vector from (0,0,0) to the point given by the pixel values. As seen in the figure, pure red, green and blue are located in the three respective corners. The points placed on the vector from (0,0,0) to (255,0,0) contain different shades of red, and the same is observed for green and blue. The vector from (0,0,0), black, to (255,255,255), white, contains the gray scale representation [24].

[4] Available at http://www.mathworks.com/matlabcentral/fileexchange/24009

The HSV colour space is a very intuitive colour representation. The basic idea is that the HSV colour representation simulates a painter mixing paint. The painter would choose a pure colour and lighten it by adding white paint or darkening it by adding black paint. The pure colour corresponds to the *hue*, the *saturation* is lowered by adding white to the pure colour and the intensity of the colour is lowered by adding black paint. The *value* then corresponds to $value = max(R, G, B)$ [24]. Figure 3.10 b) shows an illustration of the HSV colour representation.

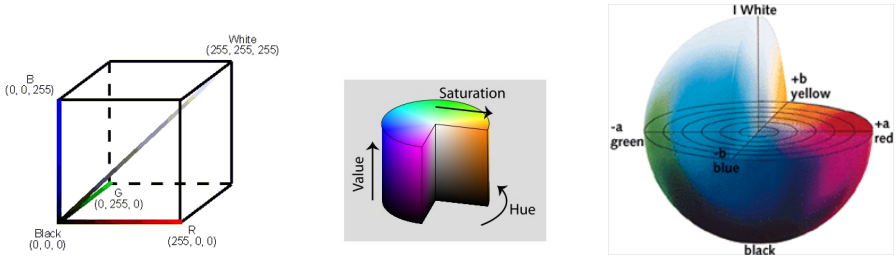

(a) RGB illustration (b) HSV illustration (c) CIELAB illustration

Figure 3.10: a) RGB colour space [5] b)Illustration of the HSV representation of colour [6] , c) CIELAB colour representation [7]

A third colour space used to quantifying iris colour is the CIELAB colour space [9]. The colour system was developed by the Commision Internationale de L'Eclairage, who wanted to make a model that was able to approximate human colour vision. The CIELAB colour space is based on three different channels, one for luminance and two for colour representation. The luminance, L^* ranges from black to white, and the two colour channels, a^* and b^* range from green to red, and blue to yellow, respectively[6]. An illustration of the colour representation can be seen in Figure 3.10 c).

3.3.2 Pixel classification

In multiple areas of this thesis a pixel classification problem occurs. One problem is the task of detecting the eyelid boundaries. The theory behind pixel classification will be explained using eyelid detection as an example.

The eyelid classification problem can be stated as: given a pixel value $d_i \in \mathbb{R}$

[5] From http://www.machinevision.ca/machinevissupport
[6] From http://www.brighthub.com/multimedia/publishing/articles/122040.aspx
[7] From http://www.sapdesignguild.org/resources/glossary_color/index1.htmlcs_lab

in the configuration of pixels, D, in an image, classify to one of the C classes, x_c, where x is a set of random variables $\{x_i, ..., x_n\}$, and $c \in \{1...C\}$ [25]. In order to classify each pixel a classification rule is needed. Often the mean and standard deviations for each class is used and a simple way to classify is to apply a threshold on these statistics. However, it is known that pixels are spatially dependent, meaning that neighbouring pixels tend to belong to the same class. We therefore want to use a method that takes the spatially dependency into account. Markov Random Fields (MRF) is such a method and will be the basis of all pixel classifications throughout this thesis.

Markov Random Fields The MRF method assumes that a pixel value can be estimated from its neighbouring pixel values. More formally this means that all the information needed to determine the class, x_i, of a specific pixel surrounded by a neighbourhood, $\mathcal{N}_i$, is found in that specific neighbourhood and is independent from the classes of the remaining pixels in the image.

However, the task is not to find the class of a single pixel, but to find the most probable combination of classes for all pixels in the image configuration, $p(x|D)$. $p(x|D)$ is called the posterior distribution and is given by Bayes Theorem [8]:

$$P(x|D) = \frac{P(D|x)P(x)}{P(D)} \tag{3.9}$$

where $P(D|x)$ is the likelihood of the configuration and $P(x)$ is the prior for the classes and $P(D)$ is the prior for the image pixels values. In the MRF framework the segmentation problem is solved by looking at energies. The Bayesian probabilities can be described in form of energies using the Gibbs distribution [3] given by

$$P(x) = \frac{1}{Z} e^{-\frac{1}{T}E(x)} \tag{3.10}$$

where T is a temperature constant that is set to 1 and Z is a normalisation factor. The probabilities from Bayes Theorem can be described in form of energies. The posterior energy is given by [8]:

$$E(x|D) = E(x) + \alpha E(D|x) \tag{3.11}$$

where the first term is a smoothing term, also called the prior energy, and the second term is a data term. The constant α is introduced as a weighting parameter of the two terms and is chosen empirically.

The prior energy for the classes is defined such that neighbouring pixels are dependent on each other. Therefore a punishment for neighbouring pixels be-

longing to different classes is introduced [18]

$$E(x) = \sum_i \left(\sum_{j \in \mathcal{N}_i} \delta(x_i, x_j) \right) \tag{3.12}$$

where $\delta(x_i, x_j)$ takes the form

$$\delta(x_i, x_j) = \begin{cases} \beta, & \text{if } x_i \neq x_j \\ 0, & \text{if } x_i = x_j \end{cases} \tag{3.13}$$

Since images often contain noise components, the pixel values may vary. However, the data for class c is generally assumed normal distributed with, $N(\mu_c, \sigma_c^2)$. Using this assumption the likelihood is given by [8]:

$$P(D|x) = \frac{1}{\sqrt{2\pi\sigma_c^2}} e^{\frac{(D_i - \mu_c)^2}{2\sigma_c^2}} \tag{3.14}$$

By taking the log of Equation 3.14, the likelihood is expressed by its energy form given by [8]:

$$E(D|x) = \sum_i \frac{\log(\sigma_c^2) + (D_i - \mu_c)\sigma_c^{-2}(D_i - \mu_c)}{2} \tag{3.15}$$

The posterior energy function can then be determined by [8]:

$$E(x|D) = E(x) + E(D|x) \tag{3.16}$$

$$= \sum_i \left(\sum_{j \in \mathcal{N}_i} \delta(x_i, x_j) + \frac{\log(\sigma_c^2) + (D_i - \mu_c)\sigma_c^{-2}(D_i - \mu_c)}{2} \right) \tag{3.17}$$

In our case the data value is the mean μ_c and standard deviation σ_c of class x_c.

Now the framework of MRF has been set up using energy functions. In order to solve the segmentation problem we are interested in minimizing the posterior energy of the configuration. One method to solve this problem is Graph Cut.

Graph Cut The Graph Cut method is an adaptation of an algorithm originally designed to find the maximum flow and minimum cut [18]. In order to perform a binary segmentation, a two class Graph Cut approach can be used. The idea is illustrated in Figure A.2. The figure illustrates a two class problem, where every pixel has a neighbourhood dependency given by the edge weight, β. Furthermore, it has a cost to be assigned to either the source or the sink. In the case of $\beta = 0$, the minimum cut would be to choose the class with the minimum cost for all pixels. Therefore raising the beta value will result in a larger spatial dependency.

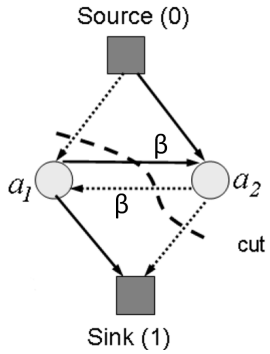

Figure 3.11: Illustration of the two label Graph Cut approach. Modified from [18]

Alpha expansion Alpha expansion is a method used for MRF segmentation of multiple classes. The basic idea of alpha expansion is to successively segment pixels into two classes, alpha and not-alpha. The alpha class changes for each iteration, and the method starts with an initial labelling. Alpha expansion uses MRF and Graph Cut in order to segment the pixels, using the constraint that pixels belonging to the alpha class cannot change class.

Training Since the human eye is very good at determining to which class a pixel belongs, the algorithm can be trained accordingly. A basic training procedure is to make human experts annotate the classes in one or multiple images. Based on the annotations, a mean and standard deviation can be calculated for each class. The classification problem can then be solved using the mean and standard deviation in the MRF energy functions.

3.3.3 Eyelid detection methods

Our method to avoid the eyelid areas, is to generate a mask for each iris map accounting for the eyelid regions. An example of an iris map and corresponding mask is seen in Figure 3.12. This section will explain how the eyelid boundaries are located and how the mask is generated.

The first step in the process of locating eyelid boundaries is to change to an appropriate colour space. The colour space should be chosen so that the difference between skin and inner eye is as large as possible. The RGB and HSV components are seen in Figure 3.13, where it is seen that the RGB colour space

Figure 3.12: Left: Extracted iris region. Right: Mask for the extracted iris region

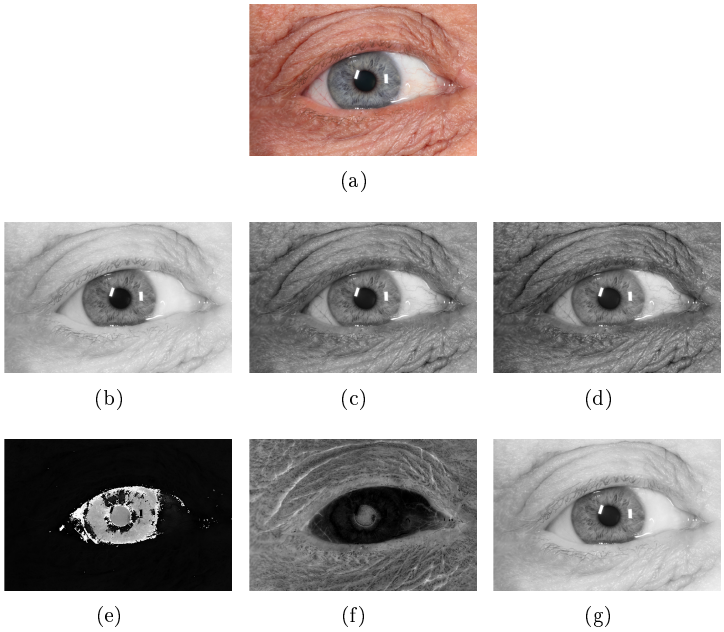

Figure 3.13: a) Eye image b) First RGB component c) Second RGB component d) Third RGB component e) First HSV component f) Second HSV component g) Third HSV component

is not very appropriate for the task. On the other hand, the second HSV component look promising, since the pixel values differ between the skin and inner eye areas, and this component is therefore chosen for further processing.

Statistics for training In the process of detecting eyelid boundaries the appropriate classes were defined as sclera, eyelashes, skin and iris. Since a variety of iris colours exist, we decided that a blue and brown class would be representative for the iris regions. In order to perform a pixel classification, there is a need for prior information about the statistics for each class. The statistics used are the mean and standard deviation for each class. The mean and standard deviation are calculated based on human annotations of the areas for sclera,

eyelashes and skin, as well as human annotations of the iris for blue and brown eyes, respectively. Ten annotations of each class were made. An example of the annotations can be seen in Figure 3.14. The mean and standard deviations, calculated using the annotations, can be seen for all classes in Table 3.1.

	Sclera	Skin	Eyelashes	Brown	Blue
Mean	0.0930	0.3566	0.4992	0.5641	0.1278
Std	0.0469	0.0756	0.0990	0.1824	0.1141

Table 3.1: Mean and standard deviation for the classes sclera, skin, eyelashes, brown and blue

Position probability matrix When looking at the eye image data, a general behaviour in location of the skin, eyelashes, sclera, pupil and iris is seen. This information is used to make a prior assumption of the positions for each class. The prior knowledge about position of the different classes was obtained from human annotations on 50 training images. An example of an annotation of an eye image is seen in Figure 3.15. The annotated areas are used in order to generate a matrix for each class containing prior probabilities based on position. The 50 images that have been annotated differ in both position and size of the iris. The center coordinates for the 50 annotated irises are seen in Figure 3.16, and it is seen that they differ as much as 175 pixels in the horizontal direction and 150 pixels in the vertical direction.

Each annotation mask is aligned with respect to the mean iris center coordinate, and the masks are scaled with respect to the mean iris radius. The alignment and scaling can be seen for an eye image in Figure 3.17.

After alignment and scaling of the annotation masks, a probability matrix is generated for all classes. This is done by summing all the annotation masks for each class, followed by a normalization. The value 0 implies zero probability

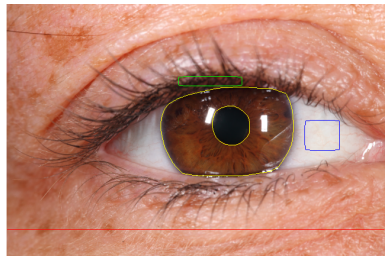

Figure 3.14: Human annotation of sclera, eyelash, skin and iris

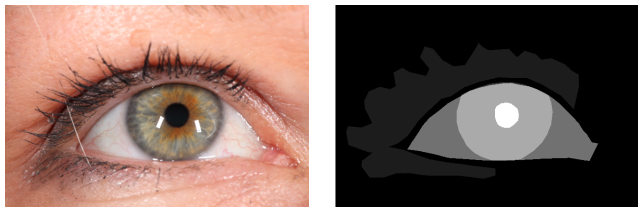

(a) Eye image (b) Annotation image

Figure 3.15: a) Example of an eye image b) Corresponding annotation of pupil(white), iris(light gray), sclera(gray) and eyelashes(dark gray). The remaining area is assumed to be skin(black).

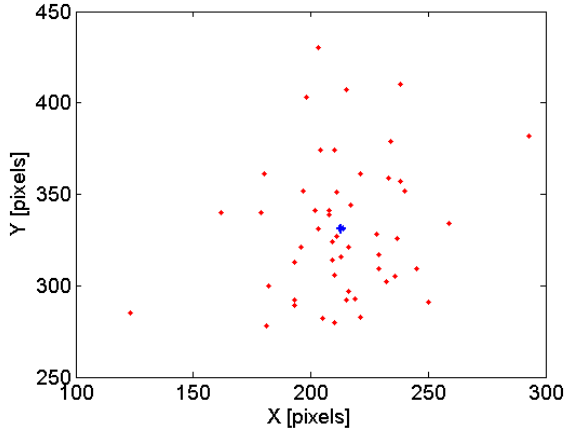

Figure 3.16: Plot of the iris center coordinates for the 50 annotated images. The blue star indicates the mean center coordinate

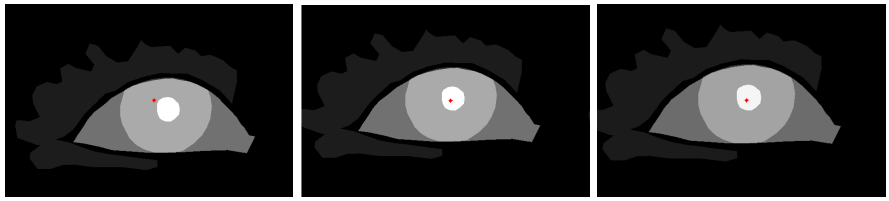

Figure 3.17: Left: Original annotation. Middle: Aligned annotation. Right: aligned and size scaled annotation. All with mean radius marked as red star

of the pixel belonging to the class and 1 implies zero probability of the pixel belonging to any other class.

The training data contains very few eye images where the upper eyelid covers a large part of the iris. The probability of having eyelid in the iris region is therefore low, leading to incorrect detection of the upper eyelid boundary. In order to account for this, we increase the probability of skin and lower the probability of iris in the joint area. The annotated skin and iris regions are therefore multiplied by 50 for a single set of annotation masks. The image was chosen due to a large occlusion of the iris. The final probability matrices are seen in Figure 3.18. The probability matrices, Q, are used in the MRF framework.

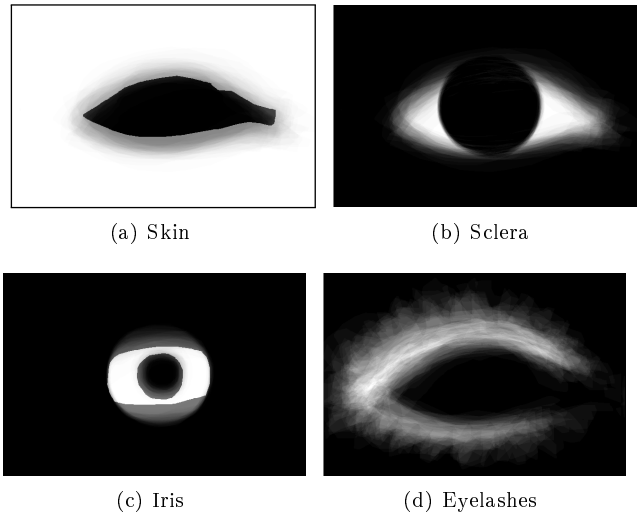

(a) Skin (b) Sclera

(c) Iris (d) Eyelashes

Figure 3.18: The figure illustrates the position probability maps for the different classes of interest

Equation 3.17 is therefore given by:

$$E(x|D) = \sum_i \left(\sum_{j \in \mathcal{N}_i} \delta(x_i, x_j) + \frac{\log(\sigma_c^2) + (D_i - \mu_c)\sigma_c^{-2}(D_i - \mu_c)}{2} \right) Q \quad (3.18)$$

The pixel classification problem using MRF is solved using Graph Cut and alpha expansion. A result is seen in Figure 3.19.

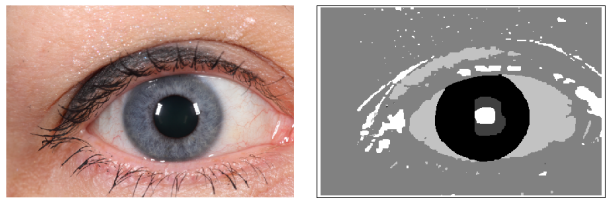

Figure 3.19: Left: Original eye image. Right: Result from MRF and alpha expansion

Fitting a spline The result from the pixel classification is a binary mask, see Figure 3.20 a). The binary mask contains values of ones in the areas classified as sclera or iris (blue or brown) and zeroes elsewhere. The mask contains noise which are seen as areas that are not part of the inner eye region. In order to remove these components, an opening procedure is applied. Furthermore, the binary mask has a hole in the inner eye region and in order to remove this, a closing procedure is applied. The two steps can be seen in Figure 3.20 b) and c).

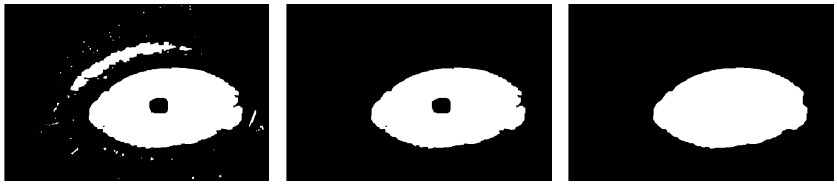

(a) Binary mask (b) Binary mask after opening procedure (c) Binary mask after closing procedure

Figure 3.20: Opening and closing procedure for binary mask

The final step in locating the eyelid boundaries is to fit a spline to the upper and lower boundary, respectively. The edges of the binary mask are found using a canny edge detector. The upper and lower eyelid regions are extracted according to the known iris center position and radius. A spline is then fitted to the upper and lower edges of the binary mask, which can be seen in Figure 3.21.

The final result for detecting the eyelid boundaries is seen in Figure 3.22.

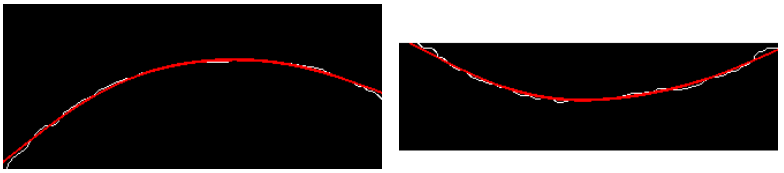

(a) Spline fitted to upper eyelid (b) Spline fitted to lower eyelid

Figure 3.21: Fitting splines. The white line is the edge and the red line is the fitted spline

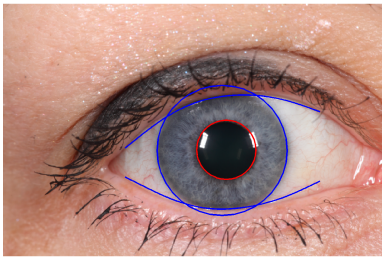

Figure 3.22: Final result for the detecting the eyelid boundaries.

3.4 Human Annotations

In order to validate our results from the iris extraction procedure a golden standard, based on human annotations of the iris boundaries, is constructed. The annotations were made by annotating along the boundaries of the pupil and iris, while making sure that the eyelids were avoided. The annotations of the iris and pupil are combined into a binary image with ones in the iris area and zeros elsewhere, an example is seen in Figure 3.23.

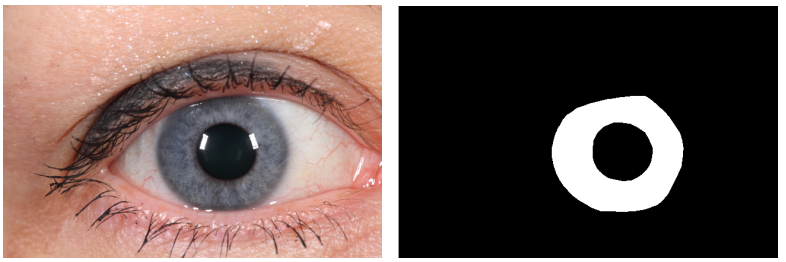

Figure 3.23: Left: Eye image. Right: Corresponding human annotation of iris area

In order to validate the results, a similarity measure, the Dice Coefficient, has been used. The Dice Coefficient is given by

$$DI = \frac{2|X \cap Y|}{|X| + |Y|} \tag{3.19}$$

where X is the detected iris area and Y is the annotated iris area. All of our eye images have been annotated manually and the Dice Coefficient has been calculated.

3.5 Iris Extraction Results

The iris extraction algorithm, explained in the previous sections, has been run on the 111 validation images. The results are compared to the golden standard which is described in Section 3.4.

The Dice Coefficient is used in order to select the images with correctly located iris regions. The iris region is assumed to be correctly located if the Dice Coefficient is above 0.92. This value was chosen from visual inspection of the results from the training images. The result for the eye image with the highest Dice Coefficient and the eye image with the lowest acceptable Dice Coefficient is seen in Figure 3.24. It is seen that the iris region is located correctly in both images. The percentage of correctly detected images, using a Dice Coefficient of 0.92 as

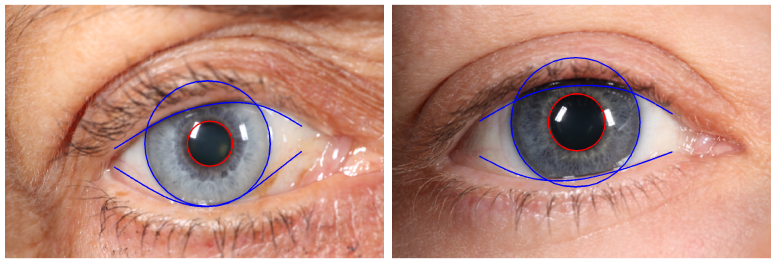

Figure 3.24: Left: Result for highest Dice coefficient. Right: Result for lowest acceptable Dice coefficient

the threshold value, is 64.3%. Six results with correctly detected iris regions are seen in Figure 3.25 and a further selection is seen in Appendix A.3. Six results with incorrectly detected iris regions are seen in Figure 3.26 and a further selection is seen in Appendix A.3. The final result does not show the performance

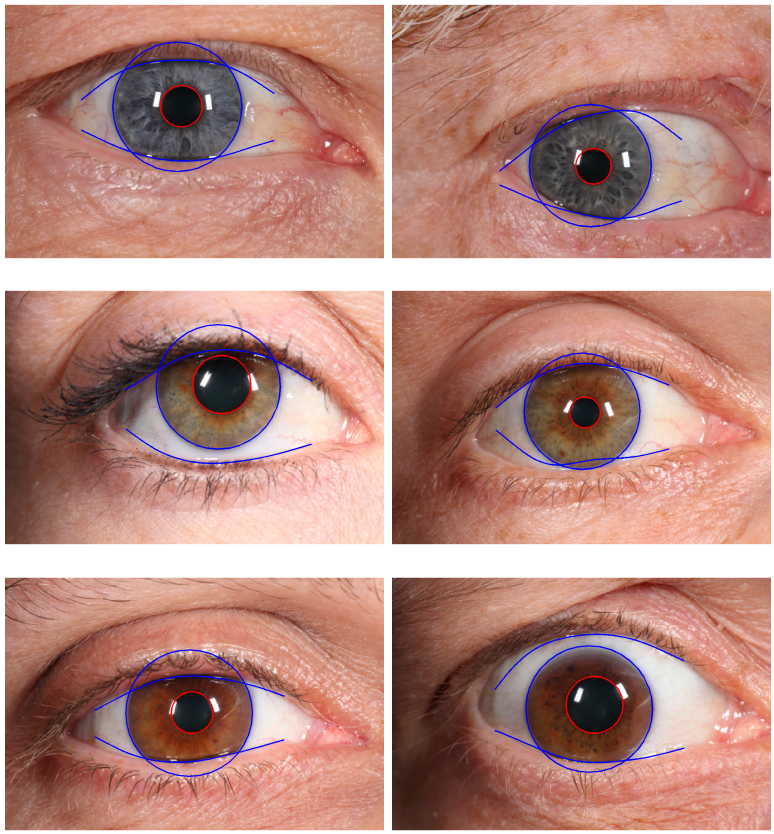

Figure 3.25: Six validation images with correctly detected iris regions

of the individual tasks. This has therefore been investigated further, and the results of the inner boundary, outer boundary, upper eyelid and lower eyelid detection is seen in Table 3.2.

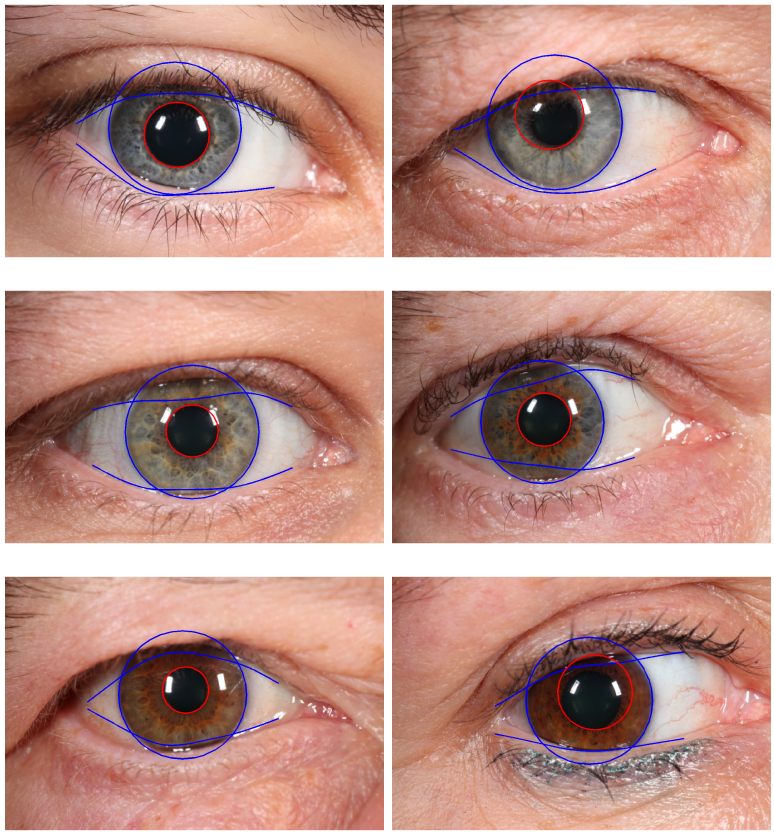

Figure 3.26: Six images with incorrectly detected iris regions

Task	Correctly detected (%)
Inner boundary	75.0
Outer boundary	96.4
Upper eyelid	76.6
Lower eyelid	73.0

Table 3.2: Individual results for the tasks performed during iris extraction

3.6 Iris Map Generation Results

The algorithm contains a radial transformation in order to transform the iris region in the original image into an iris map, see Section 3.2. The iris map has a corresponding iris mask, that accounts for the eyelid regions and reflections from the camera flashes. Iris maps and corresponding iris masks are seen in Figure 3.27 for some correctly detected images, and a further selection is seen in Appendix A.4. The iris maps and iris masks are seen in Figure 3.28 for three incorrectly detected eye images.

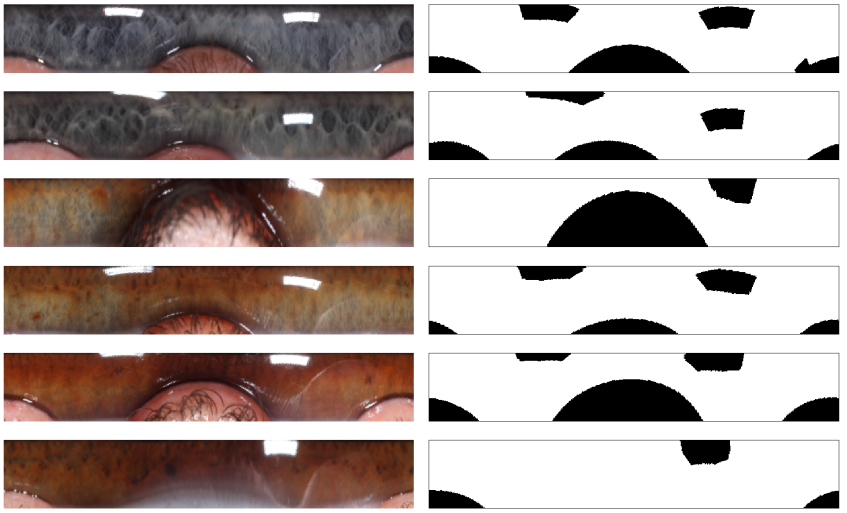

Figure 3.27: Iris maps and corresponding iris masks, for six images with correctly detected iris regions

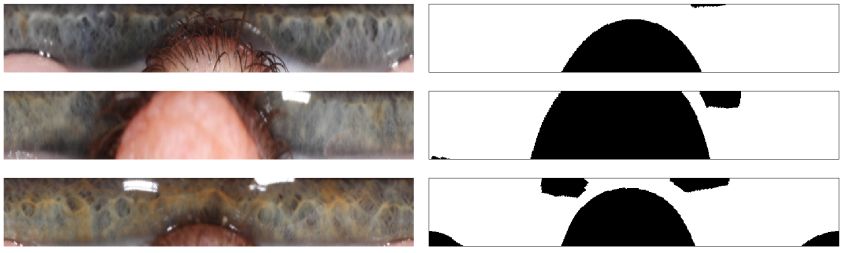

Figure 3.28: Iris maps and corresponding iris masks, for three images with incorrectly detected iris regions

3.7 Iris Extraction Discussion

Inner and outer boundary detection The individual results from the inner and outer boundary search are 75% and 96.4% correctly detected boundaries, respectively. For the inner boundary, a Dice Coefficient is calculated in order to compare the detected pupil area with the annotated pupil area. However, the Dice coefficient is inappropriate for the outer boundary, because the eyelids often occludes the iris region. The correctly detected outer boundaries are therefore identified by visual inspection.

The performance of the outer boundary detection is very high compared to the inner boundary detection. The inner boundary detection generally fails in four different cases. The first case is due to a dark iris, resulting in a low gradient between the dark coloured iris and pupil, see Figure 3.29 a). The second case arises for images with an oval pupil. In this case, the circle cannot be fitted perfectly, which leads to a incorrect detection, see Figure 3.29 b). The third case is due to the restriction that the pupil must be located inside the outer boundary. If the outer boundary is incorrectly detected and placed such that the pupil is located outside the outer boundary, the inner boundary will fail, see Figure 3.29 c). In the last case the detection will fail if a large image gradient is located as the inner boundary instead of the pupil. Figure 3.29 d) shows an eye image with a high gradient at the eyelid boundary. Furthermore the pupil boundary is occluded by eyelashes in the upper region. The highest circular gradient is therefore found on the eyelid boundary, leading to an incorrect inner boundary detection.

In order to obtain a high performance some restrictions on the size of the iris and pupil were implemented into the algorithm. The size of the iris and pupil diameters must be in the range $[11.44, 15]$ mm and $[1.44, 8]$ mm, respectively. These restrictions make the algorithm very sensitive towards the size of the eye in the image. The boundaries can therefore e.g. not be located in an image containing an entire facial region, without adjusting the parameters.

Eyelid detection The detection of eyelid boundaries is very difficult, however we succeeded in finding a method for eyelid detection. The percentage of correctly detected eyelid boundaries was 76.6% for the upper eyelid and 73.0% for the lower eyelid.

The algorithm fails in a number of different cases. An example is seen in 3.29 c). The first case is when the inner or outer boundary detection fails. This could for instance be when the outer boundary is located far from the iris region. The eyelid regions are extracted according to the known iris center position and

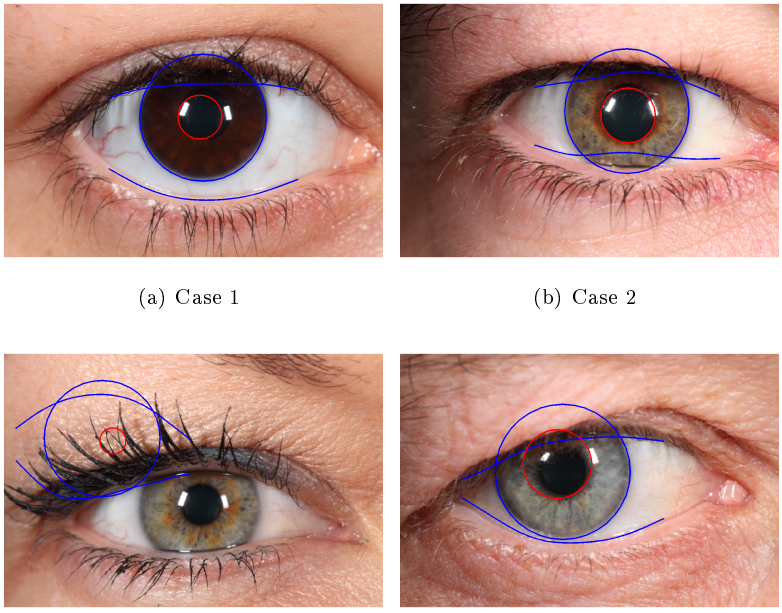

(a) Case 1 (b) Case 2

(c) Case 3 (d) Case 4

Figure 3.29: Failed inner boundary due to a) Dark iris b) Oval pupil c) Failed outer boundary d) Occlusion of boundary and high gradient at eyelid boundary

radius, and will therefore be incorrectly extracted.

The second case occurs for brown eye images. It is observed that the percentage of blue eyes among the correctly detected images are 85.9%, and the percentage of blue eyes among the incorrectly detected images are 57.5%. Therefore we can conclude that the algorithm perform better for blue eyes. The reason for the higher performance of eyelid detection for blue eye images is that the difference between the second HSV value for blue and skin is very high, see Figure 3.30 b). On the contrary the difference in the second HSV value between brown or intermediate and skin is low, see Figure 3.30 d). Obviously it is easiest to detect an area that is very different from its surroundings. In order to raise the percentage of correctly detected irises, the implementation should therefore be improved for images containing brown irises.

The third case is in the detection of eyelid boundaries for the intermediate eye images. The failure might be due to the choice of classes for the pixel

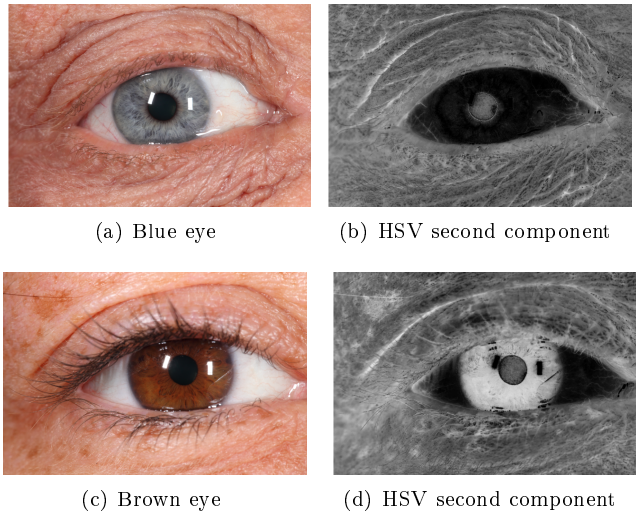

(a) Blue eye (b) HSV second component

(c) Brown eye (d) HSV second component

Figure 3.30: Second HSV component for blue and brown eye image

classification. The iris region can contain many colours from the lightest blue to the darkest brown, but in order to restrict the number of classes we chose only a blue and a brown class. The very low amount of correctly detected eyelid boundaries on intermediate eye images could indicate that more classes are necessary. The problem with the intermediate eye colours can be examined more closely. The result from the pixel classification using MRF can be seen in Figure 3.31 b). It is seen that the upper region of the iris is classified as skin instead of iris. This is due to the fact that the intermediate eye colour resembles skin more than the blue or brown class. The binary mask generated from the pixel classification contain a hole in the upper region, as seen in Figure 3.31 c). Looking at d) and e) it is seen that this results in a problem with fitting a spline to the edge of the inner eye area.

Another problem with the brown and intermediate eyes is, that they often contains darker eyelashes than the blue eyes. The dark eyelashes might obscure the boundary between iris region and skin. This gives rise to a problem since the eyelashes resemble the dark iris region, and the boundary can therefore not be located. The eyelid method is trained and optimized on images captured with specific capturing conditions. The performance of the algorithm would therefore be affected if different conditions were used. The eyelid detection is very sensitive towards changes in the colour balance of the images. This is due to the fact that the eyelid detector is trained on eye images with one specific colour balance. If eye images with either a different capturing condition or colour balance are

(a) Eye image (b) Pixel classification (c) Binary mask

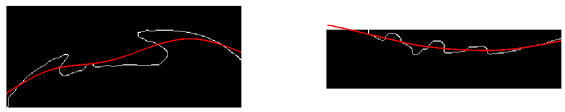

(d) Upper spline (e) Lower spline

Figure 3.31: a) Eye image b) Result from pixel classification using MRF c) Binary mask d) Fitting upper spline to edge of binary mask e) Fitting lower spline to edge of binary mask. The red circles in b) and c) indicates the problem area

used, the eyelid detector should be re-trained.

Images for further processing We have succeeded in developing a method for iris extraction. In total we obtained a correct detection percentage of 64.3%. The percentage reflects the strictness of the Dice Coefficient.

For further processing it is an advantage to include as many eye images as possible. Some of the images, which were detected incorrectly according to the chosen Dice Coefficient, can actually be used for this purpose. These images contain correctly detected inner and outer boundaries, and the eyelid boundaries are found such that the extracted iris does not contain noise from the eyelids. However, due to the incorrectly detected eyelid boundaries, the iris region for these images is reduced, see Figure 3.32. All the images that are included for further processing, even though they were classified as incorrectly detected, are seen in Figure 3.32 and 3.33.

Furthermore, some eye images, which have been correctly detected by the Dice Coefficient, are removed, since they contain small regions of skin area. The images are seen in Figure 3.34. The images are removed in order to avoid the

noise component from the skin area. The skin area will influence the statistical measures calculated from the iris region.

Modifying the pool of images for further processing, based on the above, therefore results in 83 images out of the 111 validation images. That is 74.8% of the validation images are to be used for further processing.

Figure 3.32: Incorrectly detected images, that are used for further processing

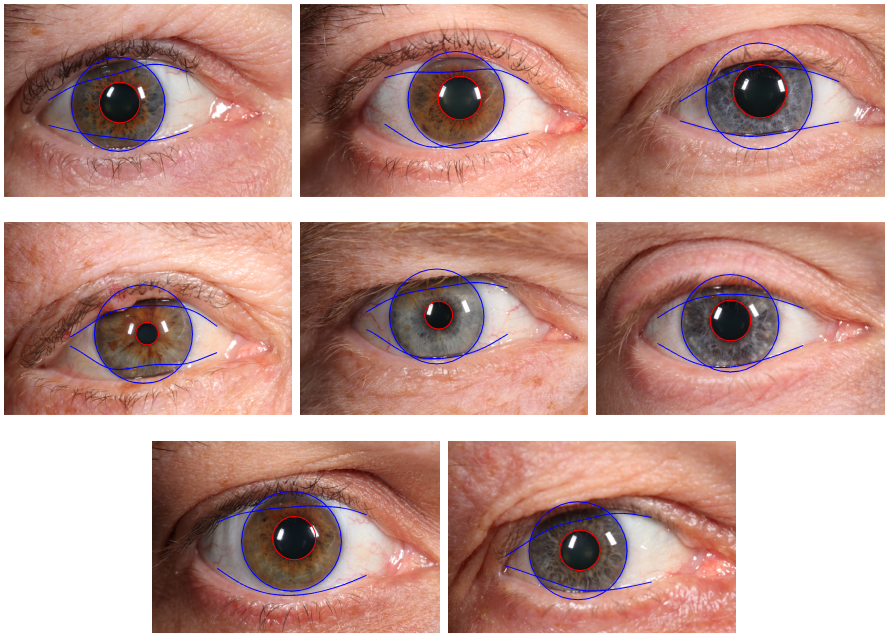

Figure 3.33: Incorrectly detected images, that are used for further processing

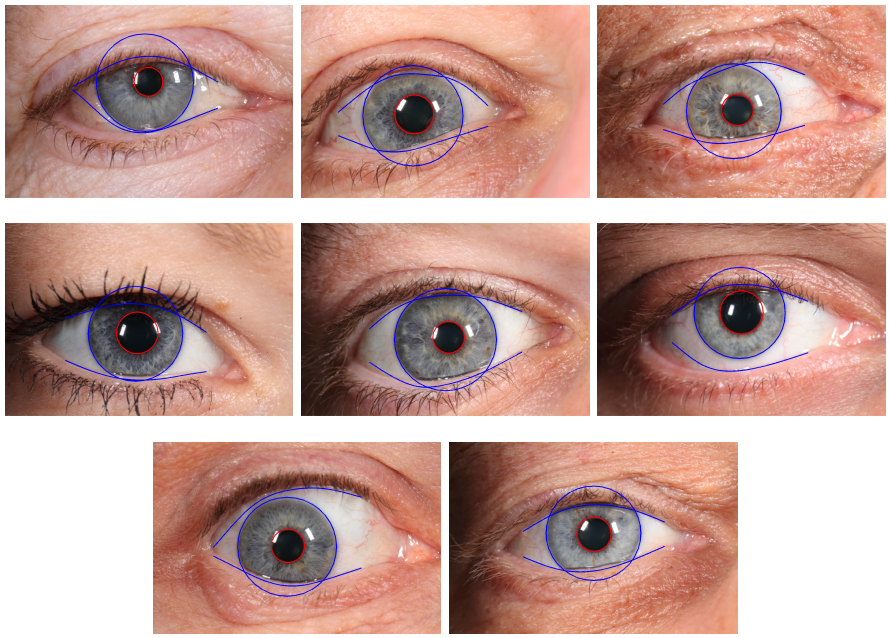

Figure 3.34: Correctly detected images, that are removed from further processing

Chapter 4

Colour Classifier

The most important feature of the iris is the mean eye colour. The mean eye colour is the only iris feature that is perceived when looking at a person face to face. Other features of the iris are only visible when looking a person deeply in the eyes or when looking at a high resolution eye image. It is therefore of interest to be able to predict eye colour from a DNA sample.

In order to predict the eye colour from a DNA sample, a quantified measure for the eye colour is needed. The simplest way to quantify eye colour is by dividing many eye images into groups of similar colour. The groups could then be correlated with a number DNA sequences of interest, in order to locate the associated DNA components. Other researchers have carried out this correlation [35, 23], by using subjective evaluations of eye colour when partitioning the eyes into groups. In this section we will explain our approach for generating an objective colour classifier, that can classify an eye image into the same class each time, thereby making the study replicable.

Two different approaches are tested for our colour classifier. In order to explain the colours in the iris map, a simple method, using statistics based on the colours in the iris map is applied along with a more advanced method using a histograms.

The colour classifier is trained, using a subjective evaluation of the eye colours,

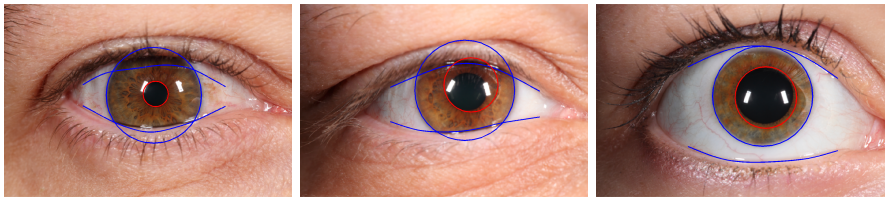

Figure 4.1: The three incorrectly detected light brown images that are included additionally

and the result from the colour classifier is tested against a subjective evaluation of another set of eye images.

4.1 Data

The data used for training the colour classifier was the training dataset explained in Chapter 2. The training dataset contains 100 eye images, but only 67 eye images were accepted from the iris extraction. Therefore only these 67 eye images are used for training the colour classifier.

Six images were subjectively classified as light brown among the 100 training images. Only one of these six images was accepted from the iris extraction. Therefore three eye images, that were not accepted from the iris extraction, are included, see Figure 4.1. Without the three additional images it would be impossible for the classifier to classify an image as light brown if using $k > 3$ nearest neighbours. The three extra images will only be included in this section of the thesis, resulting in 70 training images. The incorrect detection of the iris region is less important in this section, since it is only related to eye colour and not to the arrangement of structures in the iris region.

The training images were subjected to human evaluation of the iris colour by two individuals. The six colour classes seen in Table 4.1, were defined in collaboration with the Department of Forensic Medicine, Section of Genetics.

The human evaluation of the iris colour turned out to be very difficult, and with inspiration from [28], a number of images were used as guidelines for the evaluation. The images that were used as a guideline for separating the classes are seen in Figure 4.2. In order to be classified as "light brown", the image has to be lighter than the mean brown image. The light brown image in the upper left corner is an example. To be classified as "dark brown", the image has to

Image colour class	Numerical class
Light Blue	1
Dark Blue	2
Light Intermediate	5
Dark Intermediate	6
Light Brown	3
Dark Brown	4

Table 4.1: Colour classes

be darker than the mean brown image, and the upper right is an example. The procedure is the same for the intermediate and blue images.

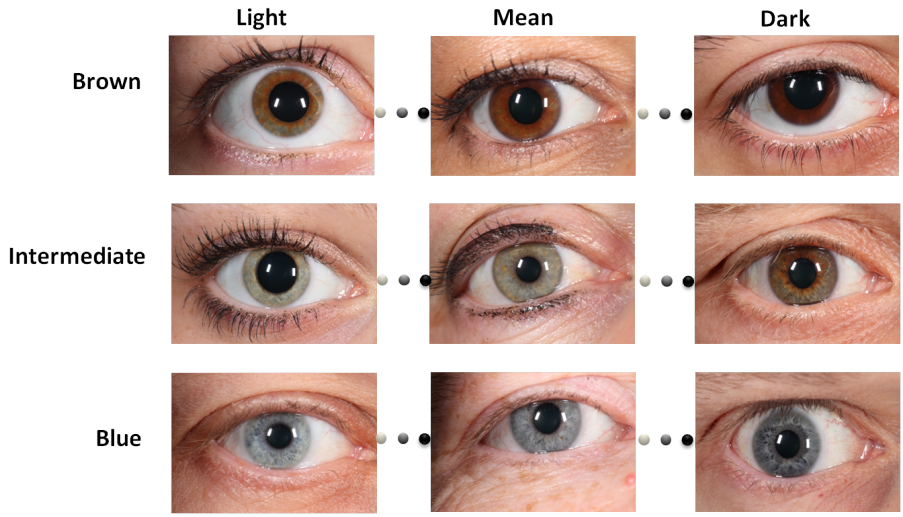

Figure 4.2: Images used for human evaluation of eye colour

The training images were classified according to the six classes by the two individuals. Even though the images were classified using the guideline images in Figure 4.2, the two classification results differed in 32% of the cases. The human evaluations for the training dataset are seen for both individuals in Appendix B.1. The approach could be to randomly choose one of the two possible classes, for every case where the classification differs. Only the evaluation from one individual was used, since choosing randomly between two evaluations would result in a non-robust classifier.

The data used for testing the colour classifier was the test dataset explained in Chapter 2. Only the 48 images that were accepted from the iris extraction were

used. The human evaluation of the images was performed in the same manner as explained for the training data.

The final results of the colour classifier were obtained by running the algorithm on the validation dataset explained in Chapter 2. Only the 83 images that were accepted from the iris extraction were used. The human evaluation of the images were performed in the same manner as explained for the training data.

4.2 Colour Classifier Methods

An objective method for evaluating the mean iris colour could be to build a classifier which classifies an image into one of the six classes given in Table 4.1. The eye image will be placed into the class most similar to the eye image, according to a given measure. Which measure to choose depends on the task to solve, but for a colour classifier the measure should consider the colours in the iris map.

Multiple measures could be appropriate for the colour classifier, but the most simple measure would be the mean RGB values of the iris map. A plot of the mean RGB values for the 70 training images is seen in Figure 4.3. One point in the figure represents one eye image, and the RGB values are the mean values calculated for all pixels in the iris map, except for the pixels avoided by the corresponding mask, as described in Chapter 3. The colour of the point represents the class that the eye image is assigned to by subjective evaluation of the eye colour, as explained in Section 4.1. The figure reveals some correlation between mean RGB values and colour class, however, there is an overlap between the different colour classes.

Section 3.3.1 explains three different colour representations, and according to the theory, the HSV space would be a more intuitive colour space. The mean HSV colour values may be the appropriate measures, and they are plotted in Figure 4.4 in the same manner as for the mean RGB values. The figure reveals that the mean HSV colour values are not an appropriate measure, since the overlap between colour classes is very high.

According to M. Edwards et. al. [9] the CIELAB is appropriate for separating between blue and brown eye colour. The mean CIELAB values are plotted in the same manner as for the mean RGB and HSV colour values, and the plot is seen in Figure 4.5. It is seen that there is a correlation between mean CIELAB values and colour class, but as for the other two colour representations there is an overlap between the classes. Investigating a simple mean value in the three

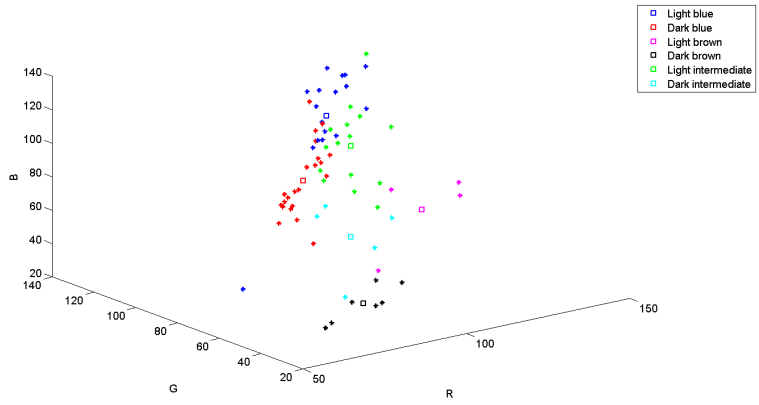

Figure 4.3: Mean RGB colours plotted for the training images

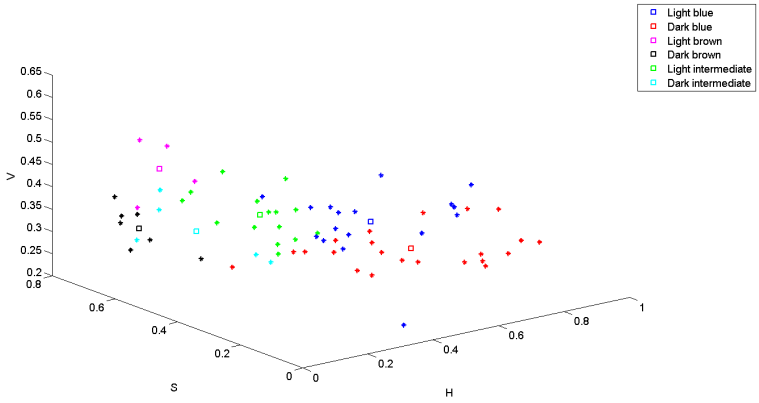

Figure 4.4: Mean HSV colours plotted for the training images

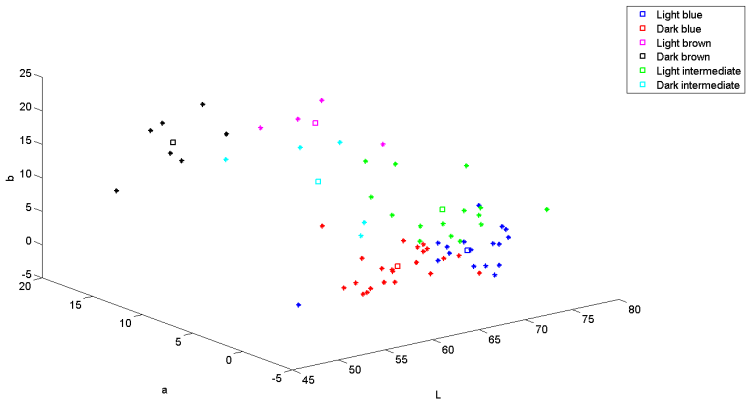

Figure 4.5: Mean CIELAB colours plotted for the training images

colour spaces RGB, HSV and CIELAB, has shown that the mean value is not appropriate as the only measure in the colour classifier. Furthermore, it was seen that the colour classes in RGB and CIELAB space have less overlap than the colour classes in HSV space.

The mean eye colour does not contain enough information about the colours represented in the iris map, and therefore a method using both mean value and coefficient of variance will be implemented for RGB and CIELAB space. Adding the coefficient of variance for each component in the colour space as a measure for the colour classifier adds more information about the colour content of the iris map.

4.2.1 Colour classifier - mean and coefficient of variance

The measures used in this colour classifier are the mean and coefficient of variance calculated for each component in the colour space. Therefore six measures are calculated for each iris map. Based on the RGB space the six measures would be mean value for R, mean value for G, mean value for B, coefficient of variance for R, coefficient of variance for G and coefficient of variance for B. The mean value μ [14], is calculated as

$$\mu = \frac{1}{N}\sum_{i=1}^{N} x_i \tag{4.1}$$

where N is the number of pixels and x_i is the i'th pixel in the iris map.

The coefficient of variance c_v [14] is calculated as

$$c_v = \frac{\sigma}{\mu} \tag{4.2}$$

where σ [14] is given by

$$\sigma = \left(\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \mu)^2\right)^{\frac{1}{2}} \tag{4.3}$$

where N is the number of pixels and x_i is the i'th pixel in the iris map

A method using k-means clustering is used in order to make the mean value μ and coefficient of variance c_v more robust. The iris map is divided into squares of 10 by 10 pixels and the mean and coefficient of variance are calculated for each area. K-means clustering is used with $k = 2$, dividing the iris map areas into two clusters. The pixels in the iris map belonging to the largest cluster is then used when calculating the final mean and coefficient of variance. The procedure is performed in order to reduce the influence of for example small brown areas in a blue eye. A small brown area in a blue iris map would change the statistical values. Using this procedure, only the most dominant eye colour is used when calculating the mean value and coefficient of variance.

The classifier is trained in the following way. The mean value μ and coefficient of variance v_c are calculated in the manner explained above for all training images. The training images are also subjectively evaluated and assigned to one of the six classes, thereby partitioning them into six groups. After training the classifier, the classifier is tested using the test dataset. A test eye image is classified by calculating the mean value and coefficient of variance and assigning it to the most occurring class of the k-nearest training images.

The distance between the test image and each training image is calculated as a sum of χ^2 distances [15]:

$$d(x,y) = \sum_i \frac{(x_i - y_i)^2}{(x_i + y_i)} \tag{4.4}$$

where x is the feature vector containing mean value and coefficient of variance for each colour component for one training image and y is the feature vector containing mean value and coefficient of variance for each colour component for the test image. In this case $i = 1...6$, because the number of measures in the feature vector is 6.

The method just explained is used for both RGB and CIELAB colour space in order to test which colour representation gives the best colour classifier.

4.2.2 Colour classifier - histogram using kd-tree

The colour classifier explained in Section 4.2.1 is only using six measures, and in this section a more complicated classifier is explained. The assumption is that more information about the iris map will result in better classification results.

The colour classifier is based on RGB space and the measure used for the comparison between training images and test image is a histogram of the iris map. A histogram of the iris map could be generated by dividing RGB space into equal sized squares and then counting the number of pixels in each square. However, such a histogram assumes that the points are uniformly distributed in RGB space. In order to conclude whether the data is uniformly distributed in RGB space, the R, G and B values from all pixels in every iris map from the training dataset are plotted. The plot is shown in Figure 4.6, where it is seen that the data is non uniformly distributed.

Since the data is non-uniformly distributed in RGB space, the histogram approach has to account for this. One method is to divide RGB space into rectangles of different size, but with an approximately equal number of pixel points in each rectangle. In order to divide RGB space into areas with an equal number of pixels, a kd-tree is used [2]. The kd-tree is generated ones and the limits for each rectangle are used to define the bins in the histogram. The resulting kd-tree is seen in Figure 4.7. As expected the rectangles are very large in the areas, that have very few pixel points and very small in areas, that contains many pixel points.

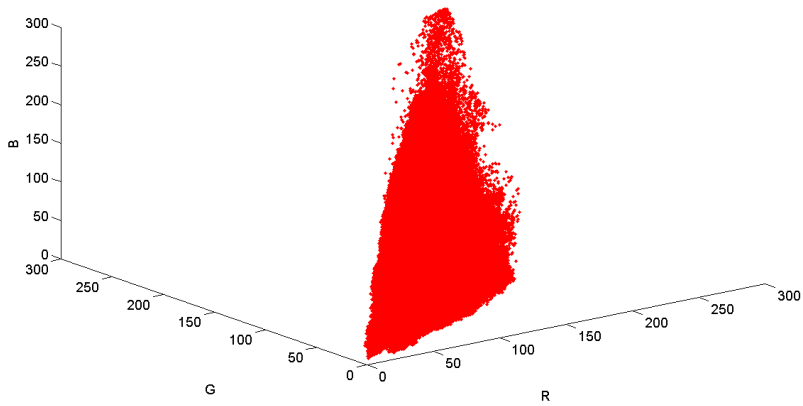

Figure 4.6: Illustration of the training data in a RGB space

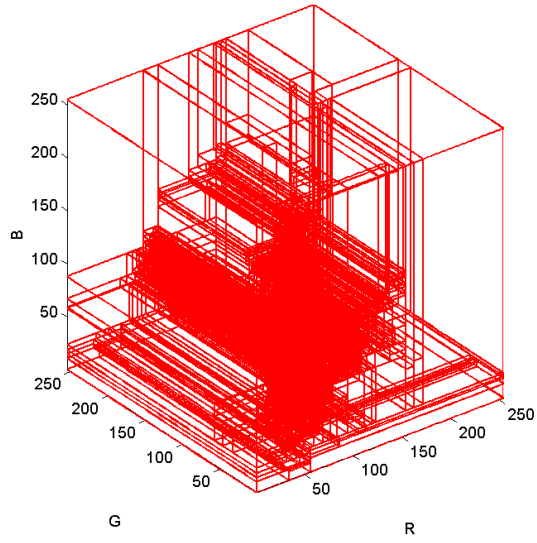

Figure 4.7: Illustration of the kd-tree of our training

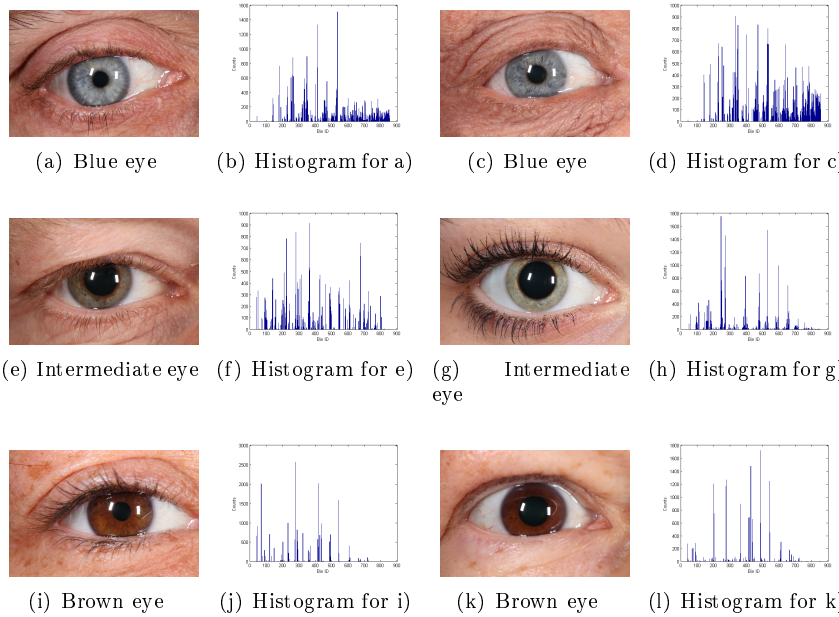

(a) Blue eye (b) Histogram for a) (c) Blue eye (d) Histogram for c)

(e) Intermediate eye (f) Histogram for e) (g) Intermediate eye (h) Histogram for g)

(i) Brown eye (j) Histogram for i) (k) Brown eye (l) Histogram for k)

Figure 4.8: Eye image and corresponding histogram for two blue eyes, two intermediate eyes and two brown eyes

A histogram is built, using the limits from the kd-tree for all eye images in the training dataset. Figure 4.8 shows two blue, two intermediate and two brown eye images and their corresponding histograms. It is very clear that the histograms differ among the three eye colours, and this indicates that the histogram could be used as measures in a colour classifier.

The colour classifier is trained by generating a histogram for all eye images in the training dataset and by subjectively evaluating the iris colour, thereby dividing the histograms into one of the six groups.

The colour classifier is tested by generating a histogram for the test image, and classifying the image according to the most occurring class of the k-nearest neighbours. The distance between two histograms is calculated as a χ^2 distance given by Equation 4.4, with $i = 1...N$, where N is the number of bins in the histogram.

4.2.3 Result of the colour classifier methods

In order to choose between the three proposed methods, the algorithms have been tested. The test images was used for this purpose. The results from testing the three approaches are seen in Table 4.2, 4.3 and 4.4. Examining the contents of the three tables reveals that the two simple approaches, using mean and coefficient of variance in RGB and CIELAB space, have lower results than the histogram approach. The highest percentage of correctly classified images is 59.8% for the simple approach in RGB space and 63.5% for the simple approach in CIELAB space. The result for the algorithm based on histograms in RGB space is 77.7% correctly classified images, for $k = 2$.

k	1	2	3	4	5
Error	25.0	23.3	21.6	20.7	19.3
Correct	23.0	24.7	26.4	27.3	28.7
Correct (%)	47.9	51.5	55.0	56.9	59.8

Table 4.2: Test of the method using mean and coefficient of variance in RGB space

k	1	2	3	4	5
Error	29	27.4	17.5	17.9	20.2
Correct	19	20.6	30.5	30.1	27.8
Correct (%)	39.6	42.9	63.5	62.7	57.9

Table 4.3: Test of the method using mean and coefficient of variance in CIELAB space

k	1	2	3	4	5
Error	13	10.7	12.2	12.7	14.2
Correct	35	37.3	35.8	35.3	33.8
Correct (%)	72.9	77.7	74.6	73.5	70.4

Table 4.4: Test of the method using a histogram based on RGB space

Looking at the results from the test of the three colour classifiers, the conclusion is that the best colour classifier is the algorithm using a histogram approach. The k-value for the k-nearest neighbour is chosen to be $k = 2$, because the highest percentage of correctly classified images.
The highest k-value tested was $k = 5$ because of the restricted number of eye images classified as light brown. Only three images classified as light brown were present in the training data, and choosing a k-value higher than 5 would result in a classifier which was not able to classify into light brown.

Final result The colour classifier using the histogram approach is chosen as the final colour classifier. The colour classifier is applied on the validation dataset, in order to get the final results. The validation images have not been used for training or testing the algorithm and is therefore completely new to the algorithm. The result is seen in Table 4.5. The percentage of correctly classified eye images is 65.4%.

Error	Correct	Correct (%)
28.7	54.3	65.4

Table 4.5: Final results from the colour classifier

4.3 Dimensionality Reduction

The final colour classifier is based on a histogram generated in RGB space, and it contains 852 bins. This is a very large number of dimensions regarding a correlation with DNA. In order to be able to correlate the histogram information with DNA, a dimensionality reduction must be performed. A very common method for dimensionality reduction is Principal Component Analysis (PCA), which is a method that chooses the components with the highest variances as the principal components [29, 13].

The idea is to perform dimensionality reduction on the histograms in order to obtain a few components that explain most of the information in the histogram. Since the histograms are a series of counts, the histograms follow a Poisson distribution. The PCA method assumes normally distributed data and in order to account for the Poisson distribution of the histograms, an Anscombe transformation is performed [22]:

$$X \to 2\sqrt{X + \frac{3}{8}} \tag{4.5}$$

The transformation into PCA space is performed using a training set of 70 histograms from our training dataset. The matrix containing the 70 histograms is the data matrix, X. The histograms in X are centred according to the mean [29, 13], μ, of the histograms:

$$X = X - \mu \tag{4.6}$$

A rotation is performed in order to change the coordinate system to a coordinate system with the first principal axis along the direction with highest variance. The second principal axis is orthogonal to the first axis and in the direction of highest variance. The rotation results in a data representation, S, given in a new coordinate system:

$$S = XL \tag{4.7}$$

where L is a linear transformation matrix [29, 13].
Figure 4.9 shows the explained variance as a function of the number of principal components. 95.58% of the variance is explained by six principal components. Therefore we can represent the information contained in a histogram by six principal components. This number of dimensions is suitable for correlation with DNA.

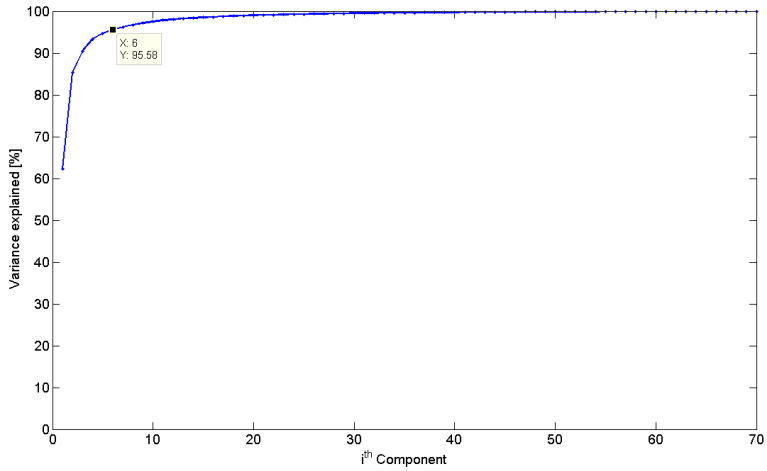

Figure 4.9: Explained variance as a function of number of principal components

In order to obtain the score for each of the six principal components for our validation dataset, the mean μ is subtracted and the linear transformation matrix L is applied on each individual histogram. The two first principal components are plotted for each iris map in the validation dataset in Figure 4.10.

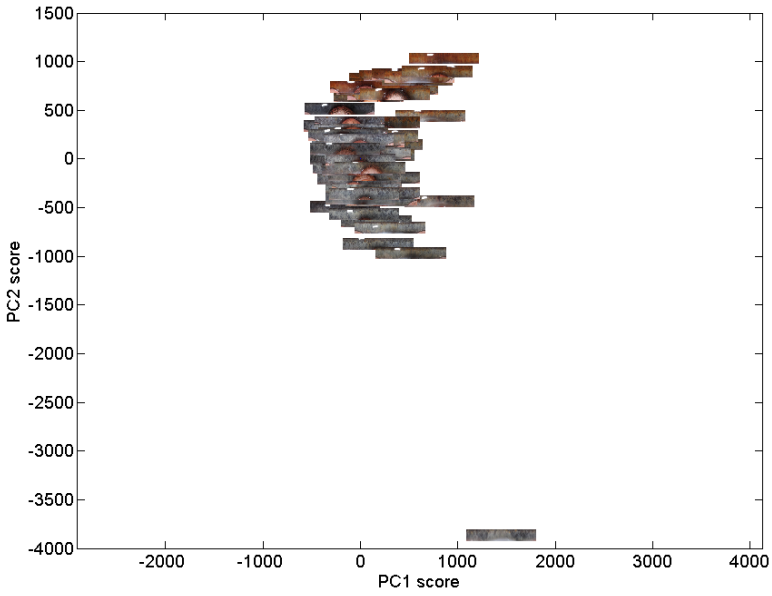

Figure 4.10: Iris maps plotted with first and second principal component

4.4 From DNA Sample to Colour Class or Synthesized Eye Colour

This section will concentrate on the possibility of creating meaningful results after a DNA correlation has been performed with the six first principal components. Since the DNA correlation has not been made yet, this section is purely hypothetical, and the results shown are based on eye images where a dimensionality reduction followed by a back transformation into a histogram has been performed. The results are therefore only generated in order to show the possibilities of the method.

Correlation between principal components and DNA can result in the ability to predict values for the principal components. The values for the six principal components can therefore be predicted for any new DNA sample. However, these six values are not very easy to interpret. The idea is therefore to assign one of the six colour classes given in Table 4.1 to a DNA sample, based on the six principal components. Furthermore, it would be very intuitive to synthesize the eye colour based on the six predicted principal components.

In order to find a colour class or synthesize an eye colour based on the six predicted principal components, a back transformation from PCA space has to be made [29, 13]. The back transformation is performed by:

$$\tilde{X} = S_i L' + \mu \tag{4.8}$$

where i is the number of principal components used. The estimated histogram, $\tilde{X}$, is therefore found by applying a reversed linear transformation of the scores for the six principal components. Furthermore, an inverse Anscombe transformation is performed, in order to account for the Poisson distribution of the histograms.

Colour class The histogram for the DNA sample has been generated by a back transformation of the six predicted principal components. This histogram can be used in the colour classifier explained in Section 4.2.2. The percentage of correctly classified eye images is 60.2%, see Table 4.6.

Error	Correct	Correct (%)
33	50	60.2

Table 4.6: Final results from the colour classifier using a histogram generated by back transformation

Synthesizing eye colour The histogram generated by back transforming the six predicted principal components can also be used for synthesizing eye colour. The approach used is simply to generate an image the size of an iris map using the histogram information. The colour chosen for a histogram bin is simply the mean colour of the RGB area in the kd-tree. The amount of each colour is then chosen based on the height of the histogram bins. This means that the number of pixels of one particular colour in the iris map is given by the percentage of counts in the given bin with respect to all counts in the histogram. The pixels are randomly positioned in the image. Examples are seen in Figure 4.11.

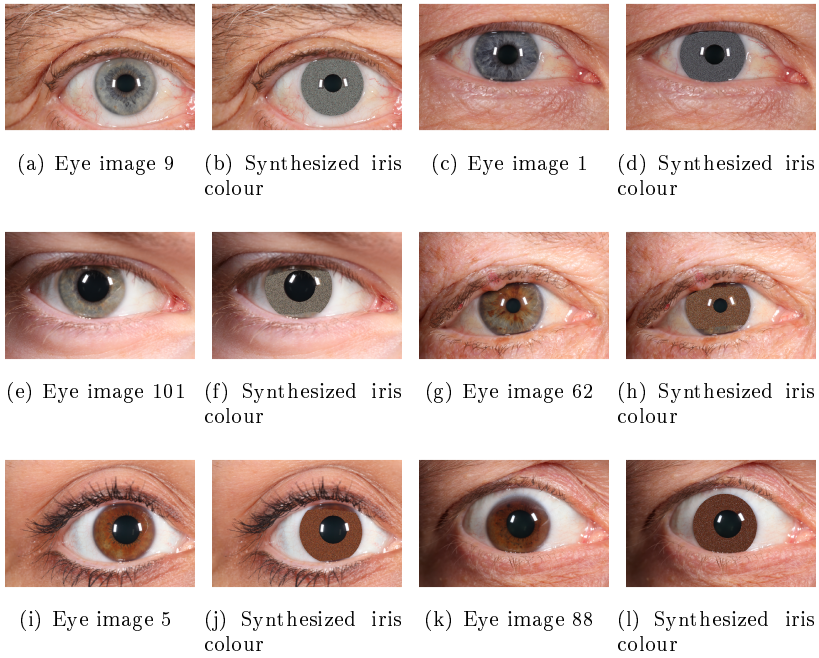

(a) Eye image 9 (b) Synthesized iris colour (c) Eye image 1 (d) Synthesized iris colour

(e) Eye image 101 (f) Synthesized iris colour (g) Eye image 62 (h) Synthesized iris colour

(i) Eye image 5 (j) Synthesized iris colour (k) Eye image 88 (l) Synthesized iris colour

Figure 4.11: Eye image and corresponding synthesized eye colour inserted in eye image

4.5 Colour Classifier Discussion

Subjective evaluation The subjective evaluation and distribution of eye images into a number of classes is a very difficult task. The difficulties arise when an eye image containing for example mainly blue colour with brown spots has to be evaluated. Does this eye image belong to the blue or intermediate class? This is a general problem when evaluating iris colour, since humans perceive colour differently.

For this thesis the subjective evaluation was performed by two individuals using the eye images in Figure 4.2 as a guideline, and the images were divided into the six classes given in Table 4.1. The percentage of disagreement between the two classification results was 32%, which is a very high percentage of disagreement under the given circumstances. The two individuals had guideline images, they had experience with looking at eye images, and they had worked together on this thesis a couple of months previous to the classification. The result strongly indicates that humans perceive eye colour differently, as was expected.

The very non-robust subjective evaluation makes it difficult to train the classifier. The classifier cannot perform better than the subjective evaluations used to train it. Our implementation is trained using only one individual's evaluation, because the inter-individual agreement was very low. The algorithm is therefore very dependent on this individual's perception of colour. This dependency could be removed by having more individuals to classify the eye images.

The subjective classification process is dependent on several considerations regarding human perception of colour, and should be thoroughly prepared. Our evaluation process was to show eye images on a computer screen. The same computer was used for all evaluations, since the colours can change between screens. The images were shown in the same order each time. This might be a problem since the perception of an iris colour can depend on the previously shown eye colours. A random order of the images could be used in order to remove the inter-image dependency. An internal control could also be applied, so that the same image would appear multiple times, thereby removing the dependence on the previous images.

Colour classifier results The result for the classifier was 77.7% of correctly classified images using the test dataset, and 65.4% when using the validation dataset. The percentage of correctly detected eye images is not impressive, and it is probably due to the difficulties with the subjective evaluation. Some examples of incorrect classification results along with the subjective evaluation are seen in Table 4.7, and the images are seen in Figure 4.12. The remaining

classification results along with the subjective evaluation are seen in Appendix B, Table B.8 and B.10. The incorrect algorithm result is often very close to the subjective evaluation, for example for image number 35, the algorithm result is light intermediate and the subjective evaluation is dark intermediate. Looking at image 35 in the figure, it can be very difficult to determine whether the colour is light or dark intermediate. Generally the algorithm fails in the cases where humans have difficulties with classifying the colour.

Image number	Subjective evaluation	Algorithm	Comparison
5	3	4	no
15	1	2	no
19	2	5	no
24	1	2	no
35	6	5	no
40	3	6	no
41	6	5	no
43	4	6	no

Table 4.7: Examples of incorrectly classified eye images

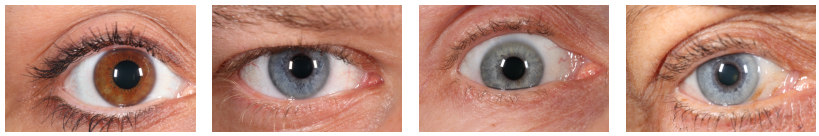

(a) Eye image 5 (b) Eye image 15 (c) Eye image 19 (d) Eye image 24

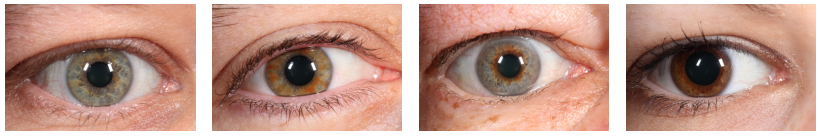

(e) Eye image 35 (f) Eye image 40 (g) Eye image 41 (h) Eye image 43

Figure 4.12: Eye images with incorrectly classified iris colour using the histogram approach

Dimensionality reduction The dimensionality reduction was performed in order to find a few explanatory variables instead of the 852 bins in the histogram. A very nice result was revealed when plotting the iris maps with respect to the first and second principal component, see Figure 4.10. The plot reveals that the second principal component seems to explain the variation between blue and brown eye colour. The histograms built for the colour classifier therefore seem to contain all the information needed to explain the colour variations. The

feature explained by the first principal component is not obvious, but might be a noise component occurring due to the capturing conditions.

Back transformation The result from the colour classification using histograms generated from a back transformation of the first six principal components is 60.2%. The percentage of correctly classified images are lowered compared to the 65.4% that was obtained for the validation dataset using normally generated histograms. This is as expected since the back transformation from principal component values to histograms introduces an error.

The synthesized iris colour images in Figure 4.11 are very equal to the actual iris colour. However, the random positioning of the different coloured pixels introduces a difference compared to the actual iris. Especially the synthesized eye images in Figure 4.11 b) and h) have a different perceived colour than the original eye images in a) and g). The difference is a result of an iris containing different coloured regions. The structures and colour distribution of the iris regions therefore seem to influence the perceived eye colour.

Chapter 5

Blue vs Brown Ratio

The colour classifier described in Chapter 4 gives an objective partitioning of eye images into six groups. This approach eliminates the problem with human perception of colour being different between subjects, and it assures that one eye image will be classified as the same class every time. In the colour classifier we have also given six continuous measures, the six principal component scores, that can be correlated with DNA. However, the six measures are difficult to interpret if no further processing is performed. Therefore a histogram is generated based on the six principal components, which can be used either in the colour classifier or to synthesize an iris colour.

Although the colour classifier solves some specific problems, a continuous measure which is more easily interpreted is of interest. Such a measure could be correlated with DNA and help with mapping the genes, coding for iris colour. The idea behind our eye colour measure is that the perceived eye colour is a result of a mixture of different shades of blue and brown. Therefore our measure was chosen to be a ratio between the amount of blue and the amount of brown in the eye image.

5.1 Blue vs Brown Ratio Method

The ratio, R_{BB}, between blue areas and brown areas in an eye image, is calculated based on the number of pixels classified as class 1, P_1, and the number of pixels classified as class 2,P_2:

$$R_{BB} = \frac{P_1 - P_2}{P_1 + P_2} \tag{5.1}$$

The ratio spans from -1 to 1, where -1 is equivalent to an image where all pixels are classified as class 2 and 1 is equivalent to an image where all pixels are classified as class 1.

Finding the ratio between blue and brown is a pixel classification problem, and the theory used for pixel classification was explained in Section 3.3.2 in connection with eyelid detection.

The first step in the pixel classification is to determine which colour space to use, and whether a single component from a colour space can be used. The pixel classification is more simple if only a single component is used. The colour space should be chosen so that the intensity difference between blue and brown is as large as possible. The RGB and HSV components are seen in Figure 5.1. None of the RGB components seems appropriate for the pixel classification, since they do not contain the wanted difference between blue and brown. The first and second component in HSV space both contain a large intensity difference between blue and brown areas. However, the first component contains similar values for both the brown area and other unwanted areas and is therefore discarded. The second HSV component is therefore chosen for the pixel classification.

A prior knowledge about the mean and standard deviations for the two classes is needed for the pixel classification. This knowledge was obtained from human annotation of blue and brown areas in two iris maps and furthermore two entire iris maps. The four iris maps are seen in Figure 5.2. All pixels, except for reflection and eyelash areas, are used for the iris maps in a) and b). For the iris maps in c) and d) all brown pixels are used in the brown statistics, and all blue pixels are used in the blue statistics. The areas are separated with a red line, and the reflection and eyelash pixels are excluded. The mean and standard deviations are calculated based on the just explained areas, and the values are seen in Table 5.1.

The mean and standard deviations are used in the MRF energy functions for each class. The energy function used to calculate the energies was defined in Section 3.3.2 by Equation 3.17.

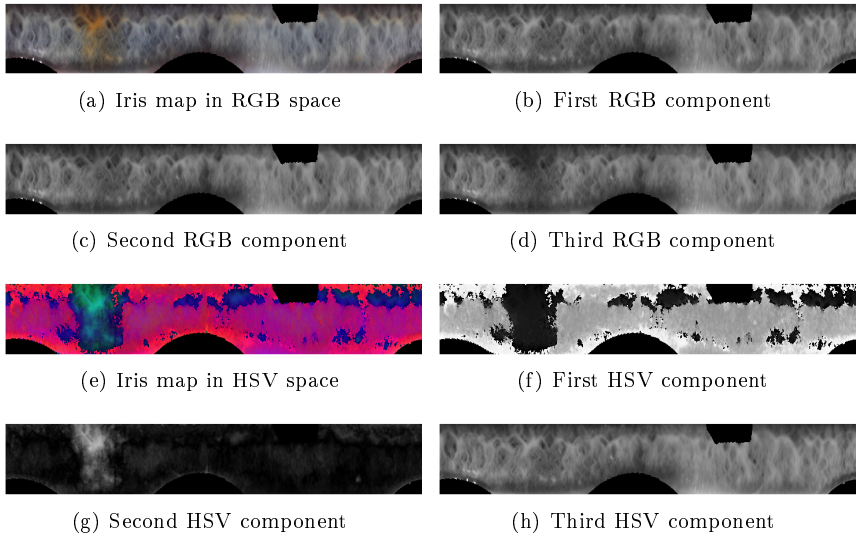

(a) Iris map in RGB space (b) First RGB component

(c) Second RGB component (d) Third RGB component

(e) Iris map in HSV space (f) First HSV component

(g) Second HSV component (h) Third HSV component

Figure 5.1: Overview of RGB and HSV colour space components for an iris map

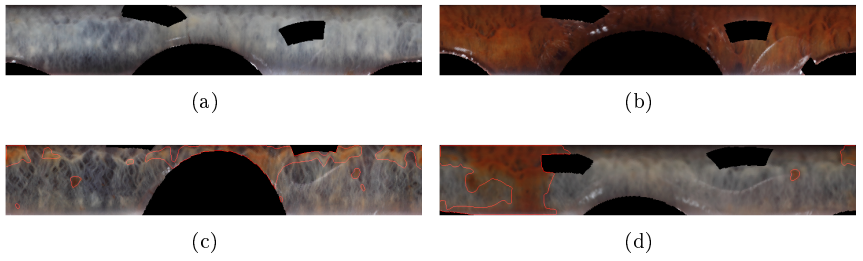

(a) (b)

(c) (d)

Figure 5.2: a) Blue area used for training b) Brown area used for training c) Brown and blue areas used for training d) Brown and blue areas used for training. Red line indicates the human annotation that separates blue and brown areas

	Brown	Blue
Mean	0.6270	0.1802
Std	0.1716	0.1983

Table 5.1: Mean and standard deviation for the classes blue and brown for the second HSV component. Calculated based on four training images

The pixel classification is performed based on the energies for the two classes and an edge weight. The edge weights are introduced in order to include a neighbourhood dependency. The edge weights are chosen to be equal for all neighbourhoods, because the dependency between image pixels are equal across an image. The neighbourhood dependency can be controlled by adjusting the edge weight, β. Figure 5.3 shows the results from MRF for three different edge weights. The results for $\beta = 0$ is a simple pixel classification without neighbourhood dependency. It is seen that single pixel areas are avoided when adding a neighbourhood dependency of $\beta = 0.1$ and that increasing the neighbourhood to $\beta = 0.3$ results in more contiguous areas. The edge weight was chosen to be $\beta = 0.1$.

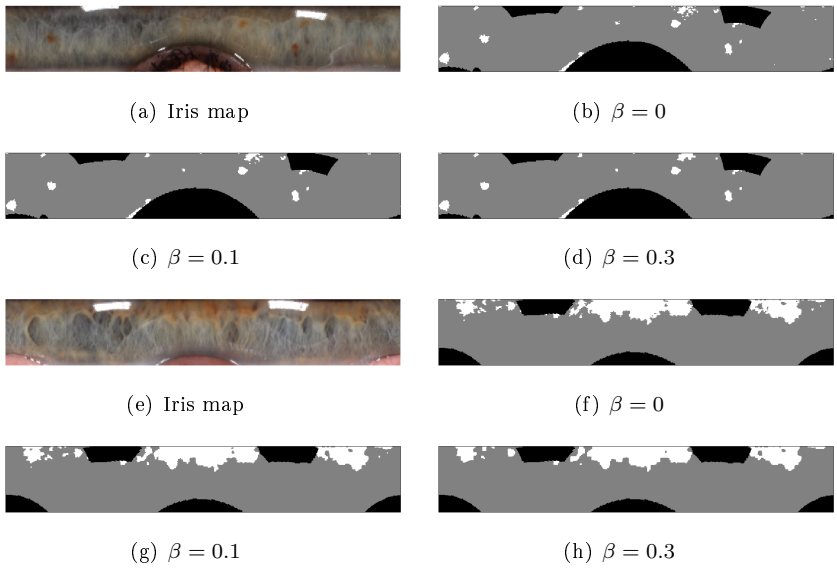

(a) Iris map (b) $\beta = 0$

(c) $\beta = 0.1$ (d) $\beta = 0.3$

(e) Iris map (f) $\beta = 0$

(g) $\beta = 0.1$ (h) $\beta = 0.3$

Figure 5.3: MRF result using different β values for two different iris maps shown in a) and e)

5.2 Blue vs Brown Ratio Results

The blue vs brown ratio is calculated for all eye images in the validation dataset, where the iris extraction is accepted. The ratios are used to generate a boxplot for each of our six colour classes, see Figure 5.4. The class for a given image is the subjectively chosen class. The red crosses in the boxplot illustrate the outliers, and the median for each class is illustrated by the red bar within the boxes.

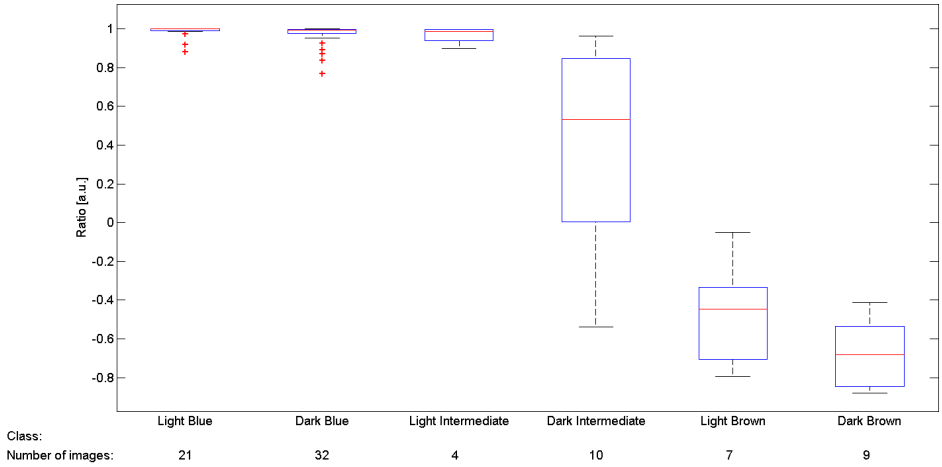

Figure 5.4: Boxplot illustrating the distribution of the blue vs brown ratio in the 83 validation images, evaluated from the subjective classification

Since the 83 validation images might not be enough data to correlate with DNA we have also calculated the ratios for the remaining images. The boxplot for these ratios can be seen in Figure 5.5.

Four examples of the blue vs brown pixel classification of the iris maps, are seen in Figure 5.6. A chosen set of extra results can be seen in Appendix C.

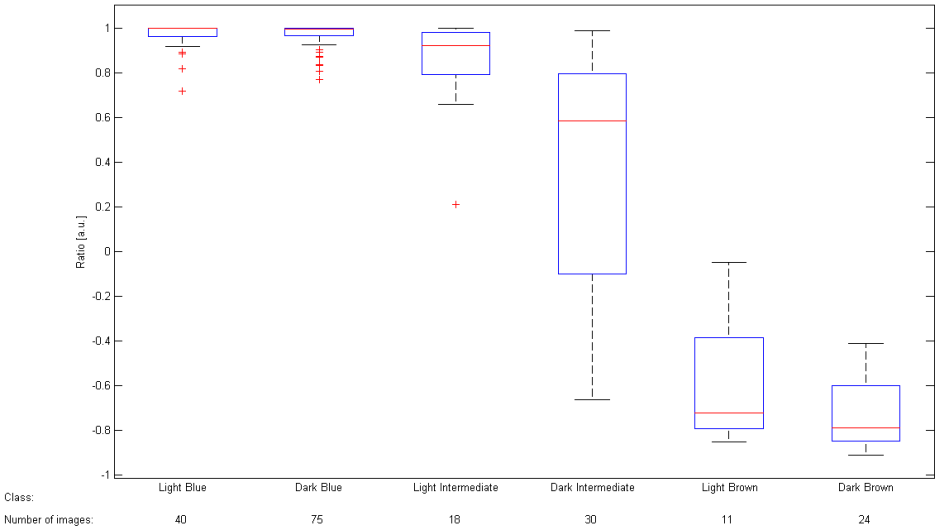

Figure 5.5: Boxplot illustrating the distribution of the blue vs brown ratio for all the eye images given for this thesis, evaluated from the subjective classification

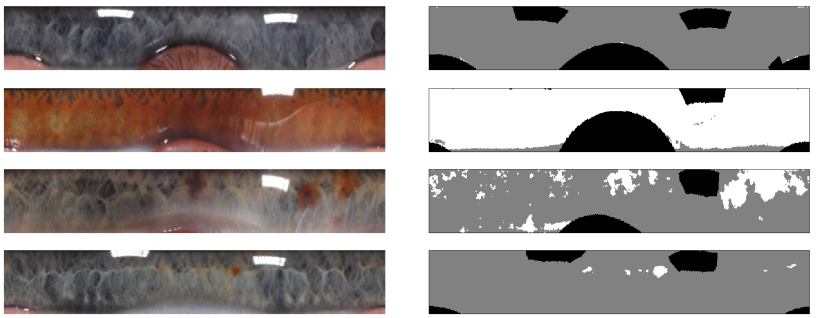

Figure 5.6: Illustrates to the left the iris map, and to the right the detected areas for different cases. White represents brown, gray represents blue and black represent areas that are avoided using the mask

5.3 Blue vs Brown Discussion

The blue vs brown ratio generated in this chapter is a very intuitive measure, since the number can be directly translated into the amount of blue and brown in the iris. The ratios have been evaluated against the subjective classification into the six colour classes given in Table 4.1. The resulting boxplot can be seen in Figure 5.4 for the 83 validation images and for all our image data in Figure 5.5.

Looking at Figure 5.5 it is seen that the mean of the two blue groups is approximately 1. The mean of the light brown group is approximately -0.7 and for the dark brown group it is -0.8. This clear separation of blue and brown based on ratios gives a promising measure for use in correlation with DNA. The intermediate groups are also separated from the blue and brown ratios based on the means. This is necessary in order to use the ratios in correlation with DNA of intermediate iris colours. The light and dark intermediate groups span the entire range from blue to brown, and the two groups can be separated based on their mean values. However, the mean of the light intermediate group is very similar to the blue groups. This either indicates that the ratio is not appropriate for separating between blue and light intermediate or that the subjective evaluation and choice of classes is inappropriate.

The boxplots in Figure 5.4 and 5.5 are very dependent on the subjective classification of the eye colours. An example where the boxplot would change appearance is if our eye images had contained very dark iris colours, as seen in for example humans of African origin. Since the eye images given for this thesis is captured in Denmark, most of the eye images contain blue, intermediate or light brown irises. This fact has influenced our perception of eye colour. Therefore the dark brown group contains brown irises that are lighter than what would probably have been the case if humans of African origin had been represented in the dataset.

The general behaviour is a smooth transition of the ratios from blue, through the intermediate, to the brown groups. This is also to be expected from the knowledge about DNA coding for eye colour, since a clear difference in the DNA between blue and brown coloured irises exists. However, the DNA coding for the intermediate groups are still to be clarified. Our hope is, that this ratio could help with identifying the genetics that result in an intermediate iris colour.

Factors influencing the ratio The ratios are calculated based on pixel classification of the iris maps. The pixel classification uses a mean and standard deviation for brown and blue. These statistics are calculated from annotated iris

areas, which have great influence on the final pixel classification and thereby the ratios. The training statistics should therefore be calculated on a representative sample of iris colour, containing different variations. Since the annotations are based on human perception of colour, the ratio is affected by our choice of training areas.

Another factor that influences the ratio is the choice of neighbourhood dependency, β, in the MRF framework. Increasing the β value introduces a higher neighbourhood dependency. The neighbourhood dependency is used in order to remove individual pixels that are falsely classified. However, too high neighbourhood dependency can result in large, interconnected and inaccurately classified areas. The choice of β value is however an estimation performed by a visual inspection of the results.

The ratio for particularly the brown eye images is influenced by large reflections present in the image. The two primary reflections from the camera flashes are accounted for in the reflection masks, but this is not possible for these large reflection areas. A brown eye image, containing such a reflection area, is seen in Figure 5.7 left. It is seen in the pixel classification result on the right, that some of the reflection area is incorrectly classified as blue. This problem affects the ratios calculated for the brown eye images, and the result is that the mean for the two brown groups are approximately -0.7 and -0.8 instead of -1 as would have been expected.

Figure 5.7: Left: iris map. Right: Result from classification. White represents brown, gray represents blue and black represent areas that are avoided using the mask

Chapter 6

Human Perception of Iris Structures

One of the main goals of this thesis is to find a method to investigate iris structures. The iris structures are not perceived when looking another human in the eye, but high resolution eye images reveal the many iris structures. It would be of great interest if a DNA correlation could provide the possibility of predicting iris structures from a DNA sample. In order to make this correlation, a quantization of the structures is necessary, and this could for instance be based on a subjective evaluation of the iris structures. However, no evaluations of human perception of iris structures has been published. In this chapter we have made an investigation of the human perception of a few interesting iris structures.

For the investigation of human perception of iris structures we have developed a GUI platform, which provides the user with the possibility to annotate some predefined structures in high resolution eye images.

6.1 Iris Structures

In corporation with the Department of Forensic Medicine, Section of Forensic Genetics, the interesting iris structures were defined as the following; Nevi dots, Fuchs Crypts, Wolfflin nodules, inner ring and outer ring. Fuchs Crypts are described as "recessed folds of the iris tissue radiating out from the pupil" [26], Wolfflin nodules are described as "lightly pigmented spots around the periphery of the iris" [26] and Nevi dots are generally defined as pigmentation placed on the surface of the iris structures [32]. A characteristic image of each of these three structure is seen in Figure 6.1.

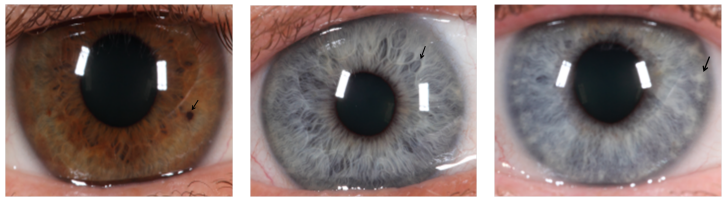

(a) Nevi dots (b) Fuchs Crypts (c) Wolfflin nodules

Figure 6.1: Interesting iris structures

The two structures, inner and outer ring, are not defined in the literature but were in collaboration with the Department of Forensic Medicine, Section of Forensic Genetics determined as interesting structures. The two structures were discovered by a thorough examination of the given eye image data. The outer ring is typically seen as a dark smooth ring placed in the periphery of the iris, and the inner ring is seen as a structural pattern that distinguish the outer part of the iris from the inner part. Characteristic images of each of the inner and outer ring are seen in Figure 6.2.

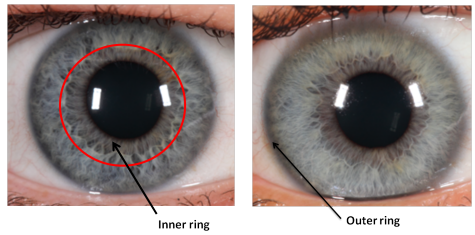

(a) Inner ring (b) Outer ring

Figure 6.2: Interesting iris structures

6.2 The GUI Platform

The GUI platform has been designed so that the user can evaluate both colour and structure of a given eye image. The eye image is loaded by pressing the load button in the main window. This displays the image and creates a pop-up window where the general colour of the iris should be chosen (not illustrated). The user is asked to choose a zoom region surrounding the iris, by clicking on the eye image. The zoom region is shown in the main window. The load and zoom process is illustrated in Figure 6.3.

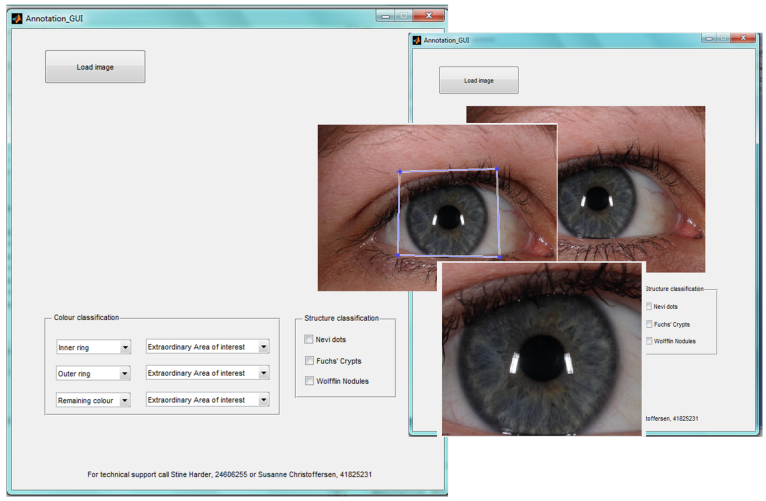

Figure 6.3: Main window for the GUI platform along with the first zooming steps that the user should perform

In the lower part of the main window, the interesting structures and colours for evaluation can be chosen. From the left, the inner ring, outer ring or extraordinary area can be chosen. The extraordinary area is used for eye images with e.g. a large brown interconnected area in a blue eye. When one of the boxes on the left are chosen, a drop down menu is activated where the user selects a colour. Afterwards the annotation tool is activated, and it is possible for the user to annotate. The process is illustrated in Figure 6.4, where the outer ring is annotated in the illustrated image. The annotated mask for the outer ring is illustrated in the lower part of the figure. A more specified guide describing the procedure step by step was provided to the seven individuals that performed the evaluation and it can be found at CD-ROM provided with this thesis.

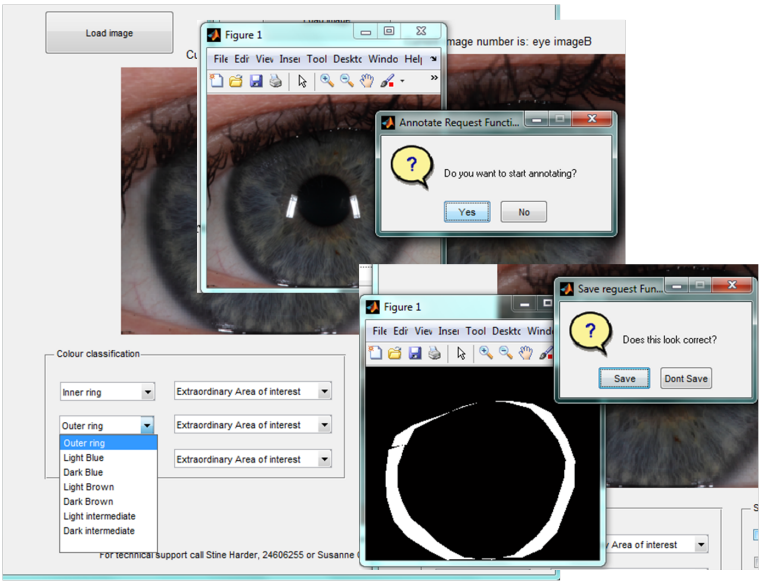

Figure 6.4: Main window for the GUI platform showing the annotation procedure for an outer ring

6.3 Human Perception of Iris Structures Method

We had seven individuals to evaluate the 111 validation images, each person locating and annotating the different structures perceived in each eye. Since the annotation process is very time consuming, the annotation process for one individual was performed over several days. All seven individuals received instruction to the different structures and to the GUI platform before they started annotating. Furthermore, all individuals were provided with the guide for the GUI.

Operator variance The most obvious measure to describe the agreement between individual operators would be to calculate the overall percentage of agreement or the effective percentage of agreement, where at least one individual identified the structure. However, these percentages do not reflect the agreement that would have appeared by chance, which would sometimes be the case [30]. The Kappa measure, which is a reliable measure for agreement is therefore used instead, since it "indicates the proportion of agreement beyond that expected by chance" [30]. The Kappa measure takes the form

$$\kappa = \frac{P_o - P_c}{1 - P_c} \tag{6.1}$$

where P_o is the proportion of observed agreement and P_c is the proportion of agreement expected by chance [30]. The proportion of agreement between the operators is given by [11]:

$$P_o = \frac{1}{Nn(n-1)} \left(\sum_{i=1}^{N} \sum_{j=1}^{k} n_{ij}^2 - Nn \right) \tag{6.2}$$

and the proportion of agreement due to chance is given by [11]:

$$P_c = \sum_{j=1}^{k} \frac{1}{Nn} \sum_{i=1}^{N} n_{ij} \tag{6.3}$$

where N is the number of images in our case, n_{ij} is the number of raters who assigned the i'th image to the j'th category, and k represents the number of categories, in our case two (present and not present) [11]. A Kappa value of 0 indicates that the agreement is no better than what would be expected by chance or guessing, and a value of 1 indicates a total agreement between the raters [30].

6.4 Human Perception of Iris Structures Results

The Kappa values calculated based on the evaluations from the seven individuals can be seen in Table 6.1. From the table it is seen that the Kappa values for the different structures are generally below 0.5. The only structure with a Kappa value above 0.5 is the Nevi dots with a Kappa value of 0.56.

Structure	Kappa
Inner Ring	0.36
Outer Ring	0.23
Nevi Dots	0.56
Fuchs Crypts	0.42
Wolfflin nodules	0.43

Table 6.1: Kappa values

6.5 Human Perception of Iris Structures Discussion

The result from the investigation of operator variance was that the Kappa values generally was below 0.5, indicating that the probability of agreement is very low. The Kappa value for the outer ring is 0.23, which is very close to the probability of agreement due to chance (Kappa = 0). However, the Nevi dots have a Kappa value of 0.56, and the probability of agreement is therefore higher than expected by chance, which is not a convincing result.

The investigation of the human perception of the different iris structures reveals that the structures with the lowest Kappa values are the inner and outer ring. The reason why these two structures are more difficult to locate than the remaining structures could be that the structures are more loosely defined. The inner and outer ring are structures determined in collaboration with the Department of Forensic Medicine, Section of Forensic Genetics, and have not previously been reported in literature. This difference, compared to the other three structures, could have resulted in a less clear definition, since we had no previous knowledge of the two structures. Fuchs Crypts and Wolfflin nodules have a similar low Kappa value, which reveals that the human perception of the structures differ between individuals. Nevi dots are, as stated earlier, the structure with the highest Kappa value. The Nevi dots differ both in colour and in structure compared to the remaining iris, especially in blue eye images where the brown Nevi dots are in contrast to the blue iris. This might be the reason why Nevi dots are the structure with the highest Kappa value, and therefore the highest inter-operator agreement.

Our result of investigating human perception of iris structure strongly indicates that humans perceive the chosen iris structures very differently. Therefore it would be impossible to build a structure classifier based on a subjective evaluation of iris structures.

Chapter 7

Image Clustering

Much research on the subject of iris colour evaluation is based on subjective evaluations, e.g. [35, 17]. The subjective evaluation makes it difficult to replicate the results, since the evaluations might change between experts and even between evaluations of the same expert. Furthermore, the subjective evaluations can result in for example a non-robust classifier, as was seen in Chapter 4. An objective evaluation of iris colour is therefore needed.

The iris structures are very difficult to evaluate by visual inspection, as concluded in Chapter 6, and furthermore, an objective method for evaluation is more appropriate. The dependency between iris colour and structure is furthermore of interest. A totally data driven approach for clustering eye images based on structure and colour is described in this chapter. The method is to make a global image descriptor based on structure and perform clustering based on this descriptor. Since iris colour is also of interest, the method will also be expanded to include colour.

A global description of iris images is wanted in order to correlate the iris structures and thereby identify common characteristics among the different irises in our database. We look at the problem as a feature matching task that has to be solved by developing a local descriptor for the structures. The approach has its origin in the *bag of words* model, which is often used when performing object recognition.

The bag of words model is based on building a dictionary of words from a set of documents [31]. This way of thinking is transferred into an image aspect by considering images with different features as the set of documents and words in the dictionary, respectively. The bag of words model for images is illustrated in the top part of Figure 7.1. The dictionary is built by clustering the features into different groups, where the center of the groups are used as a series of visual words. The occurrence of each visual word in an image is placed in a histogram that thereby describes the image information.

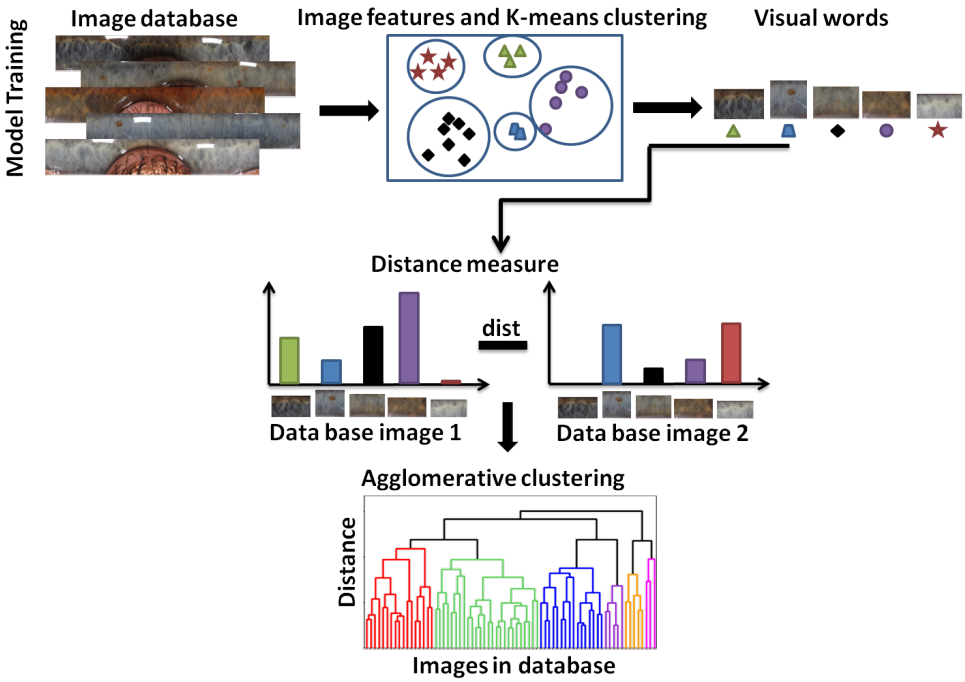

Figure 7.1: Illustration of the bag of words model for images along with the images clustering procedure

The dictionary, that is built, can be used for image comparison or clustering, as illustrated in the lower part of Figure 7.1. For each image the occurrence of each feature in the dictionary is found and placed in a histogram describing the image. The clustering of the images can be performed by calculating the distances between the histograms. Many different clustering methods can be appropriate for the task.

In order to obtain the visual features, a local image descriptor is needed. Different local image descriptors exist, but the most commonly used one is the Scale

Invariant Feature Transform SIFT descriptor [21]. The SIFT descriptor describes local image features using gradient images. From the gradient images a feature histogram is built. The SIFT method results in a local feature descriptor that is invariant to image scaling, translation and rotation. Another local image descriptor is the DAISY descriptor, which in a similar way to SIFT is based on spatial sampling of gradients. The DAISY descriptor has been shown to be much faster than the SIFT descriptor, while having the same high performance [34]. We have therefore chosen the DAISY descriptor.

7.1 DAISY Descriptor

A DAISY descriptor is a local image descriptor [33, 34], and it works in a way that is similar to for example SIFT [21]. In this thesis the DAISY descriptor is used in order to find similarities between images, and more specific it will be used to explain structures of the iris maps.

The DAISY descriptor is seen in Figure 7.2. The basic idea is to generate an eight bin histogram for each circle, describing the gradient content of the area inside the circle and then concatenate these histograms into one large descriptor of the image region. The DAISY descriptor is generated on the basis of gradient images. The gradient images in the positive x and positive y directions are calculated by running a simple filter across the image. The filter is seen in Figure 7.3. The filter basically corresponds to calculate the intensity difference between two neighbouring pixels. Once the gradient images have been found in the positive x and positive y directions, these can be used to calculate gradient images for each quantized direction:

$$I_\theta = I_x \cdot cos(\theta) + I_y \cdot sin(\theta) \tag{7.1}$$

The gradient images are according to [33, 34] called *orientation maps*, and the size of the image gradient at the location (u, v) for direction o is given by $G_o(u, v)$. The orientation maps are convolved with several Gaussian kernels, in order to obtain the circular areas. The size of the circular area is proportional to the standard deviation, Σ, of the kernel, and the circular area is weighted such that a gradient norm has less weight if it is far from the center of the circle. The orientation maps are convolved with a number of Gaussian kernels with different Σ, in order to obtain all the circular areas, and these maps are called *convolved orientation maps* [33, 34].

For every circle in Figure 7.2 an eight bin histogram is generated. The histogram vector $h_\Sigma(u, v)$ contains the values of the convolved orientation maps, using

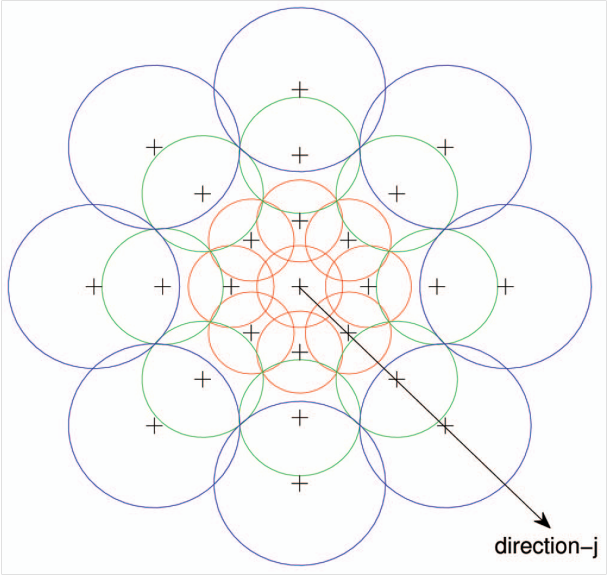

Figure 7.2: The DAISY descriptor: Each circle represents a region where the radius is proportional to the standard deviations of the Gaussian kernels, and the "+" sign represents the locations where the convolved orientation maps are sampled. From [34]

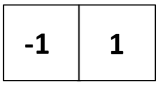

Figure 7.3: Gradient filters for finding gradients in the x direction

standard deviation Σ, at position (u, v):

$$h_\Sigma(u, v) = [G_1^\Sigma(u, v), ..., G_H^\Sigma(u, v)]^T \tag{7.2}$$

where $G_o^\Sigma(u, v)$ is the orientation map $G_o(u, v)$ convolved with a Gaussian kernel with standard deviation Σ, and H is the number of quantized directions [33, 34]. Therefore $h_\Sigma(u, v)$ is an eight bin histogram, containing one bin for each direction.

The last step is to concatenate the eight bin histograms, for all the circles, into one descriptor for the image region. The full DAISY descriptor can be defined

as:

$$
\begin{aligned}
&D(u_0, v_0) = \\
&[\tilde{h}^T_{\Sigma_1}(u_0, v_0), \\
&\tilde{h}^T_{\Sigma_1}(l_1(u_0, v_0, R_1)), ..., \tilde{h}^T_{\Sigma_1}(l_T(u_0, v_0, R_1)), \\
&\tilde{h}^T_{\Sigma_2}(l_1(u_0, v_0, R_2)), ..., \tilde{h}^T_{\Sigma_2}(l_T(u_0, v_0, R_2)), \\
&... \\
&\tilde{h}^T_{\Sigma_Q}(l_1(u_0, v_0, R_Q)), ..., \tilde{h}^T_{\Sigma_Q}(l_T(u_0, v_0, R_Q))]^T
\end{aligned}
\tag{7.3}
$$

where $l_j(u, v, R)$ is the location with distance R from (u, v) in the direction given by j, and Q is the number of circles [33, 34].

Since the number of circles, Q, is 25, see Figure 7.2, and the histograms to concatenate have eight bins, the final image descriptor is a 200 bin histogram.

7.2 Image Clustering Method

The method used to perform clustering of the iris maps is based on the bag of words model, using the DAISY descriptor as explained in the previous section. Furthermore, the method will be extended to include colour. The method consists of a training part, where the visual bag of words is build, a test part, where the visual bag of words is used to construct a histogram for each iris map, and a clustering part, where the iris maps are clustered, using hierarchical agglomerative clustering.

In order to create a bag of visual words, a set of training images have to be chosen. The training images should be representative of the entire dataset of images, and therefore a set of images containing both brown, intermediate and blue eye colour is chosen. Furthermore, the training images are chosen so that they contain structures representing the different variation in the irises. The training set can be seen in Figure 7.4.

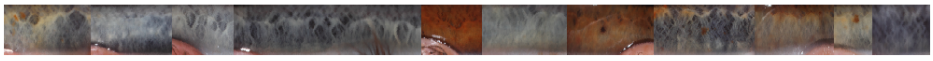

Figure 7.4: Training image

The training image is used to create a bag of visual words. The first step is to generate a large number of histograms for the training images using the DAISY descriptor explained in Section 7.1. The implementation calculates gradient images for four quantized directions, where the four directions used are the positive and negative x and y directions. The modification from Tola et al. [34] was made since the results from using four directions were similar to the results using eight directions. The histograms that are built for each circle in Figure 7.1 contain four bins, since only four directions are used. The final descriptor is $4 \cdot 25 = 100$ bins long.

The training process uses 40000 DAISY descriptors calculated at randomly chosen pixel positions. The descriptors are clustered, using a k-mean clustering. The number of clusters is chosen to be 400, which constitutes the number of visual words. The actual visual words are then the center of each cluster.

The bag of visual words is now used for generating an explanatory histogram for multiple images. For every pixel in a test image the DAISY descriptor is applied, and the generated daisy descriptor is compared to the center of each of the 400 clusters. The pixel is assigned to the nearest of the 400 cluster centres using an L2-norm given by [15]:

$$d = \sum_i (x_i - y_i)^2 \tag{7.4}$$

where i is the bin in the DAISY descriptor, x_i is the i'th bin in the center cluster and y_i is the i'th bin in the DAISY descriptor.

Figure 7.5 shows the training image where all pixels have been assigned to a visual word, that is one of the cluster centres.

The final image descriptor is then an explanatory histogram containing the occurrence of each visual word in the image. In order to maintain some of the spatial information in the images, the explanatory histogram is built using the Gaussian weighting seen in Figure 7.6. For each of the Gaussian weights, a weighted occurrence of each visual word is found, leading to an explanatory histogram of size $400 \times 12 = 4800$ bins. An explanatory histogram can be seen in Figure 7.7.

In order to match images based on local structures, an explanatory histogram has been calculated for the validation image data. A hierarchical agglomerative (HA) clustering has been used to create groups of common structures [19]. The distance metric used is the χ^2 distance given by Equation 4.4. At the first level the HA clustering finds the two closest explanatory histograms and connects them. Then the clustering continues to connect explanatory histograms to either the closest group of histograms or to the closest individual histogram. The

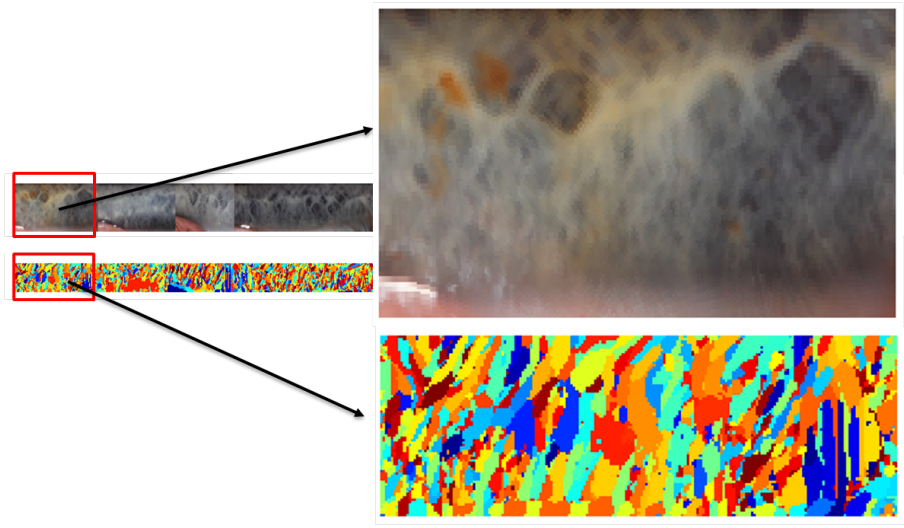

Figure 7.5: The figure shows a zoomed version of the training image indicated by the red rectangular box. The assignment of each pixel to a visual word is illustrated in the bottom image

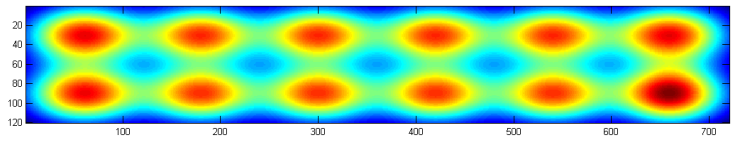

Figure 7.6: Gaussian weight

distance to a group of histograms is calculated as the distance to an average histogram of the group. In this way a tree structure is built, from where it is possible to divide the dataset into a number of clusters representing the data. An example of a cluster tree, also called a dendrogram, is seen in Figure 7.1.

The clustering has until now only been based on the local DAISY descriptor, and therefore only on the structures in the iris region. However, a data driven clustering based on colour is also requested, and furthermore some of the structures might be colour dependent. A method combining structure and colour could therefore be appropriate. For example the Nevi dots or Fuchs Crypts seems to be colour dependent, since they can be difficult to locate on an eye

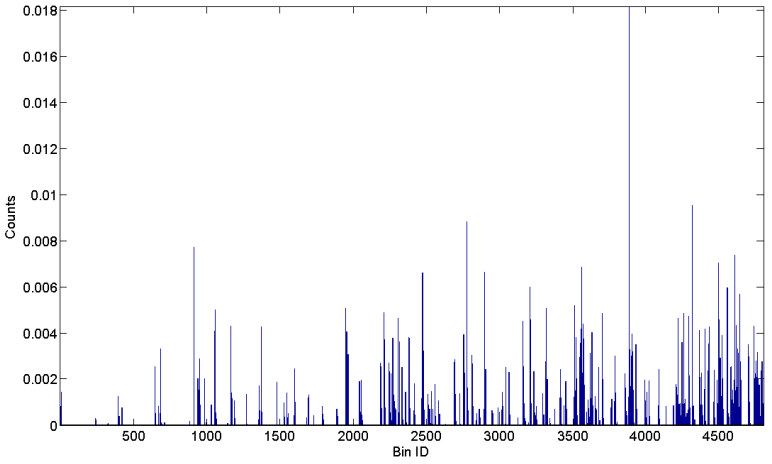

Figure 7.7: Example of an explanatory histogram for an image

image converted to gray scale, as seen in Figure 7.8. The gray scale iris map shows the difficulty with differentiating between the Nevi dots and other dark iris regions when colour is absent.

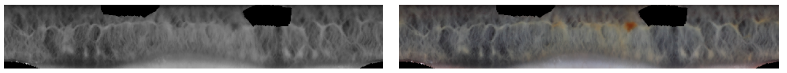

Figure 7.8: Example showing the colour dependence of Nevi dots. Left: Iris map converted to gray scale. Right: Original iris map

7.2.1 Adding colour

In order to add colour to the image clustering, a vector containing the R, G and B value is concatenated with the daisy descriptor. This is performed for each pixel in the image, and the entire colour information is therefore used. The size of the daisy descriptor with colour is 103, and the new descriptors including colour are used for building the explanatory histogram. Colour is of course added to the daisy descriptors for both training and test images. The further

process is equivalent to the clustering without colour.

The clustering method has to be evaluated, and an appropriate number of clusters have to be chosen. It is of interest to find a few large groups of images, since this is appropriate in connection with DNA correlation. However, these groups should be on such a level that the images in one cluster group have the same general tendency with respect to colour, structure or a combination of the two.

7.3 Image Clustering Results

The image clustering process is a data driven approach, however the clustering that is performed has to be evaluated in order to determine the appropriate number of cluster groups. The clustering can be seen in Figure 7.9 to 7.11 for structure, colour, and a combination of structure and colour.

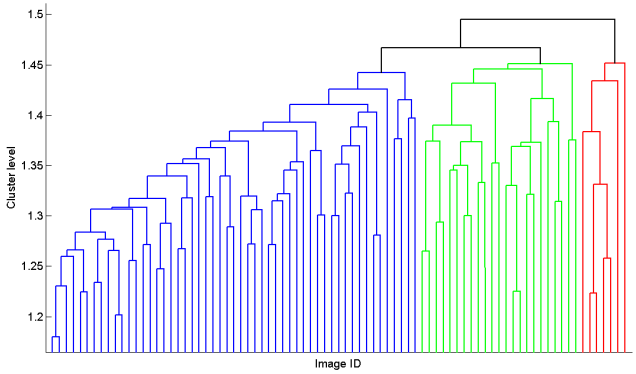

Figure 7.9: Dendrogram for structure using DAISY descriptor

In Figure 7.9 it is seen that the dendrogram for structure has three large clusters. The number of clusters was chosen with respect to the dendrogram behaviour. The iris maps for the groups are seen in Figure 7.12 to 7.15. It is seen that most of the iris maps in the first group have a high level of structure and that the second and third group have less structure. The second group also contains large areas of reflections. However, there is some images that do not seem to belong to the group in which they have been placed. The images in group 1 contain a high amount of structure, but ten images are very smooth compared

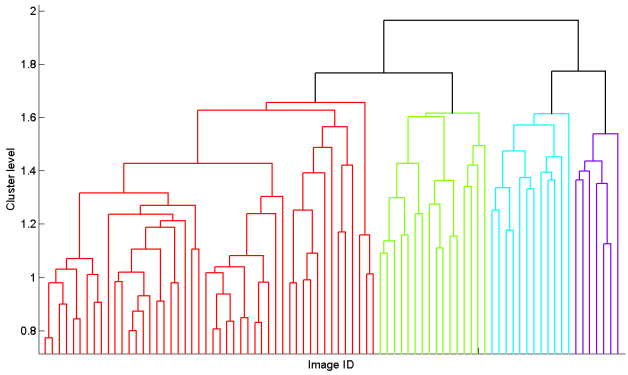

Figure 7.10: Dendrogram for colour using R, G and B values

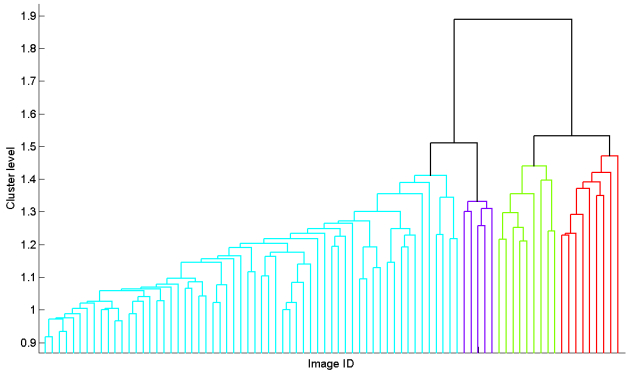

Figure 7.11: Dendrogram for structure and colour combined, using both DAISY descriptor and R, G and B values

to the remaining images. Furthermore, two images placed in group 3, have much structure compared to the remaining in this group. In order to be able to evaluate the dependency between structure and colour the structure groups are illustrated in colour in Appendix D.

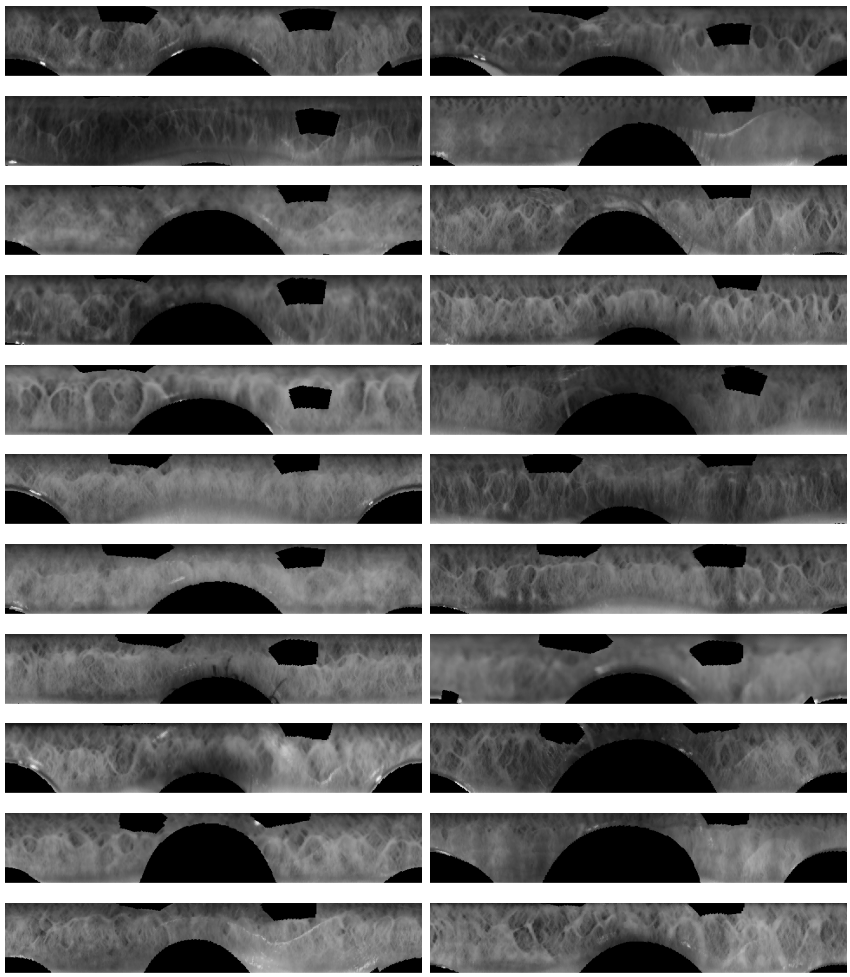

Figure 7.12: Validation images clustered based on structure - group 1

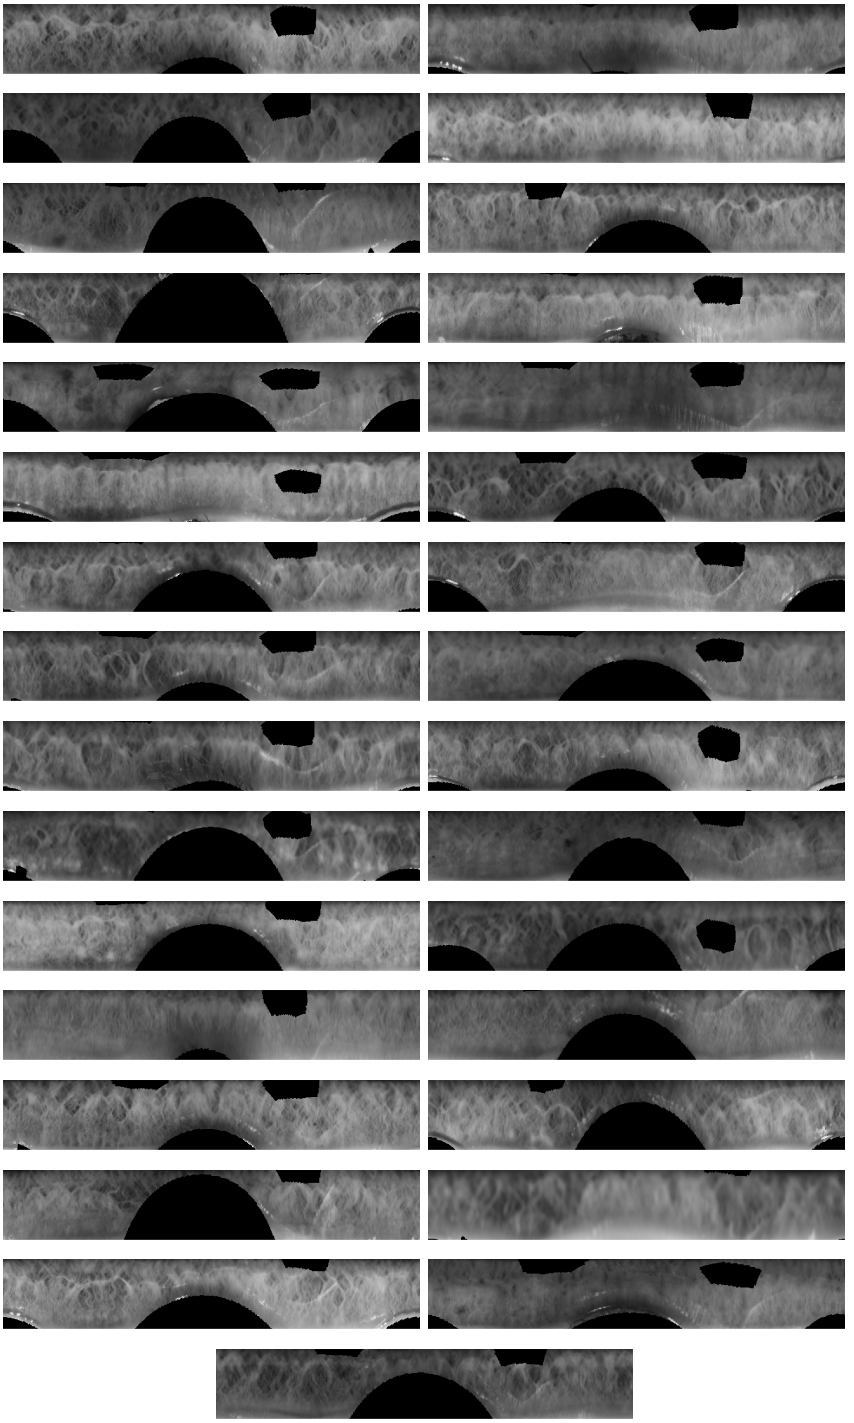

Figure 7.13: Validation images clustered based on structure - group 1

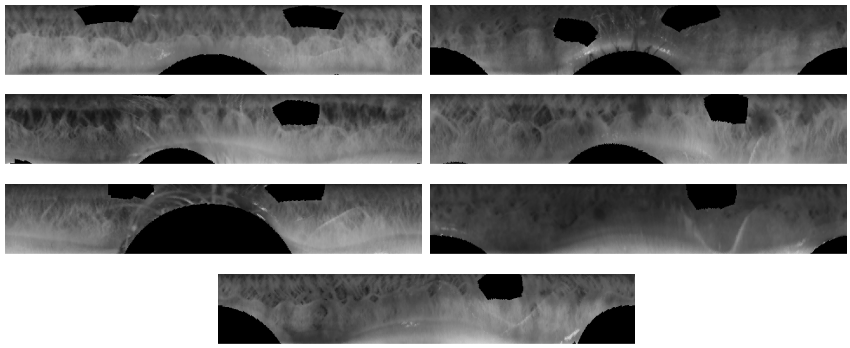

Figure 7.14: Validation images clustered based on structure - group 2

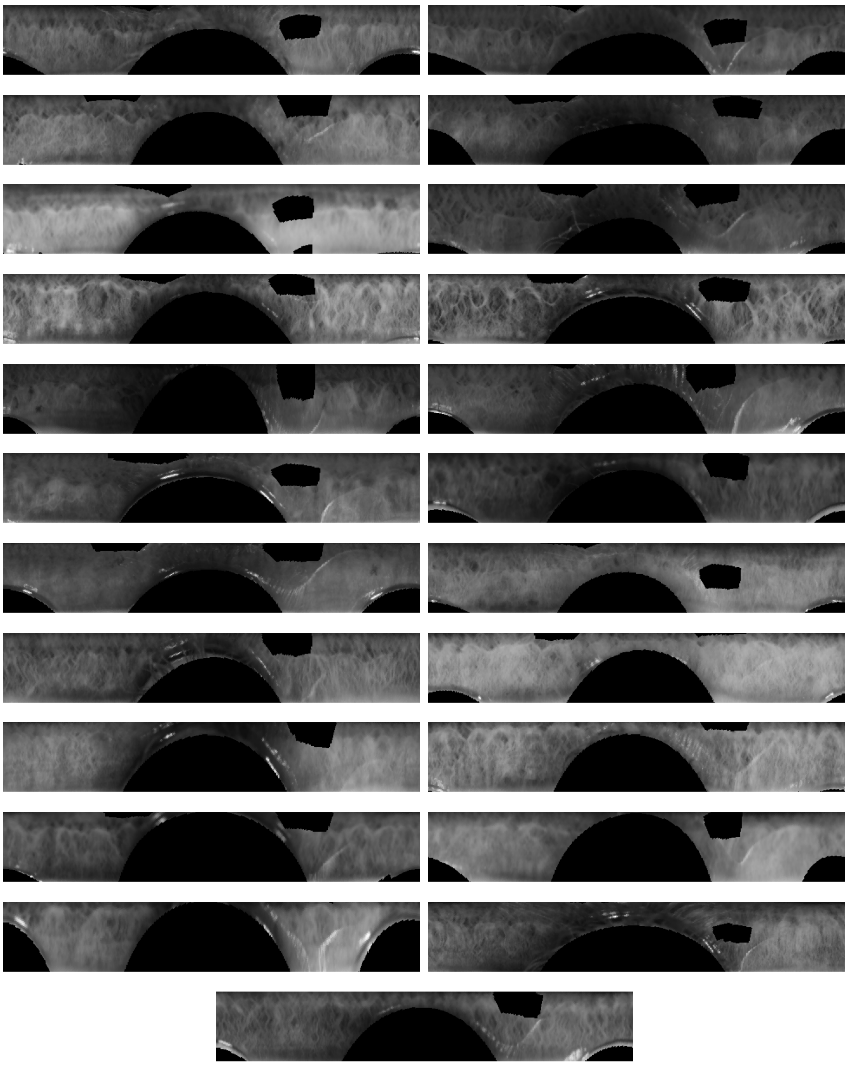

Figure 7.15: Validation images clustered based on structure - group 3

In Figure 7.10 it is seen that the dendrogram for colour has four large clusters. The four groups seems to be the most natural choice by looking at the dendrogram. The iris maps for the groups are seen in Figure 7.16 to 7.20. It is seen that a clear division of the colour has been performed. The first group contains blue eyes, the second group contains an intermediate eye colour that is mostly blue, the third group contains intermediate eye colour that is more brown than the second group and the fourth group contains brown eye colour.

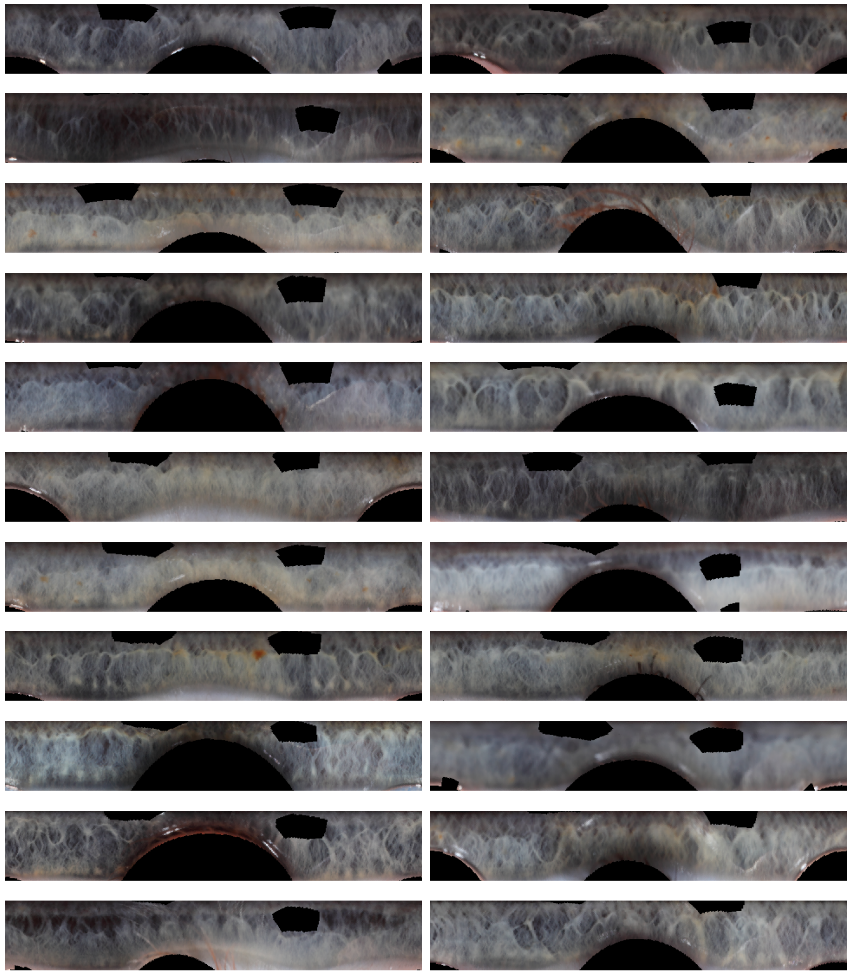

Figure 7.16: Validation images clustered based on colour - group 1

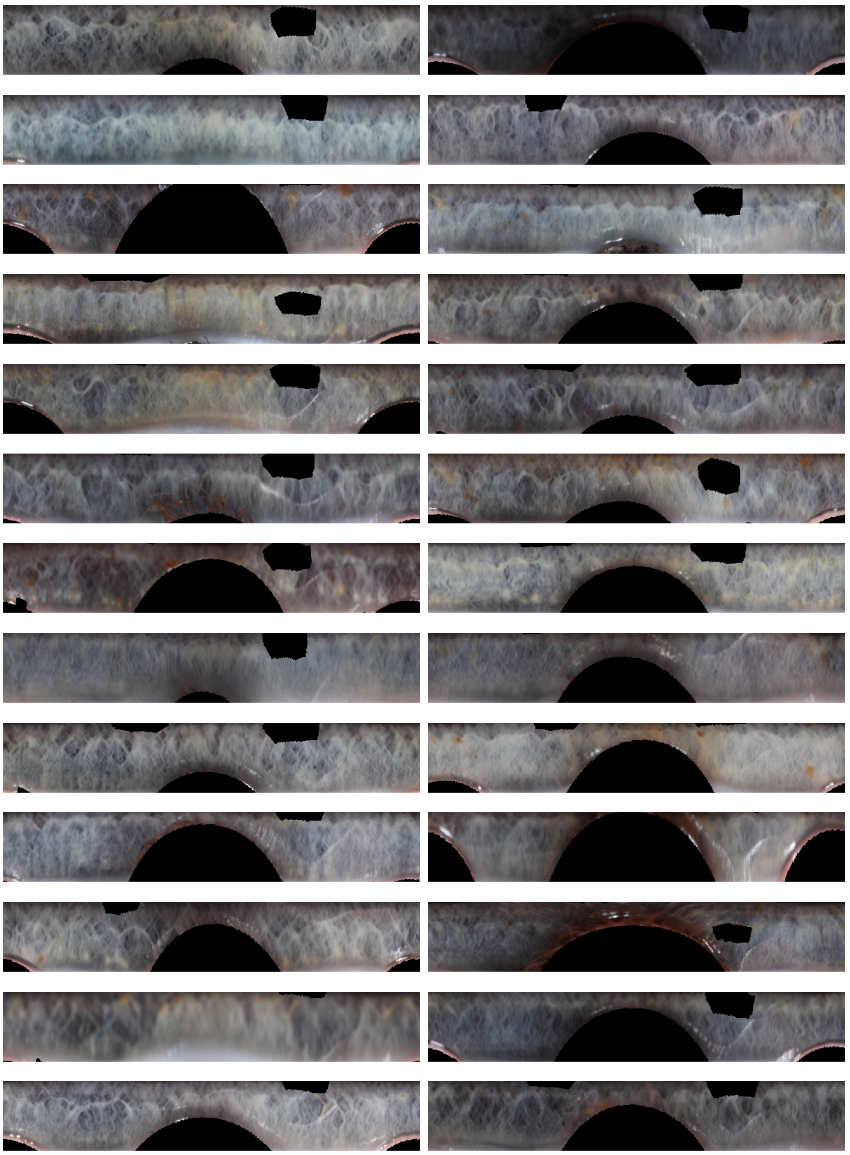

Figure 7.17: Validation images clustered based on colour - group 1

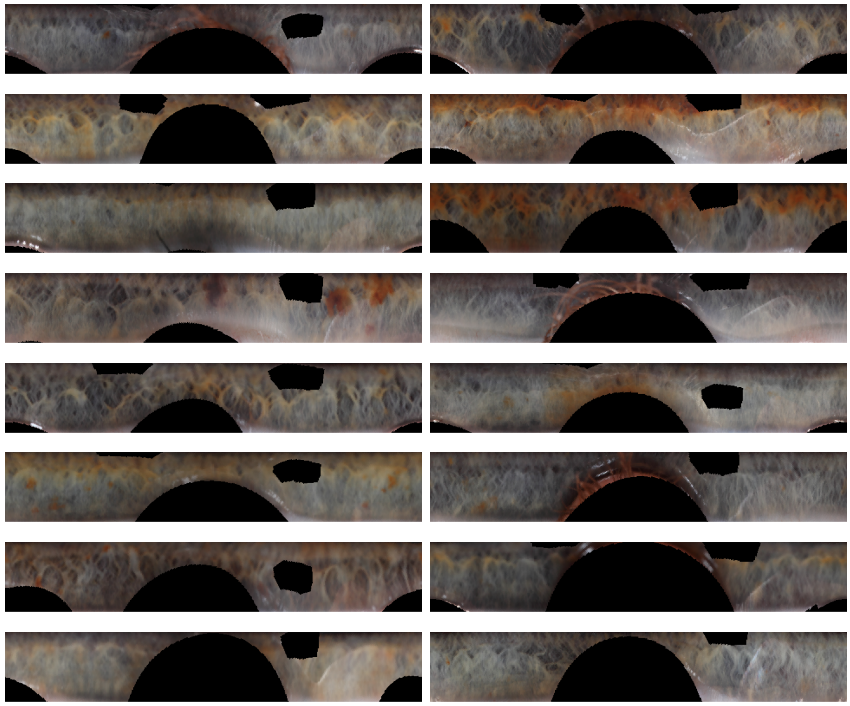

Figure 7.18: Validation images clustered based on colour - group 2

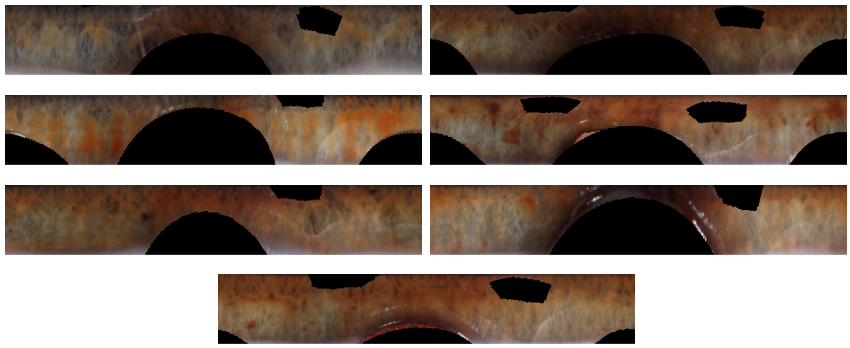

Figure 7.19: Validation images clustered based on colour - group 3

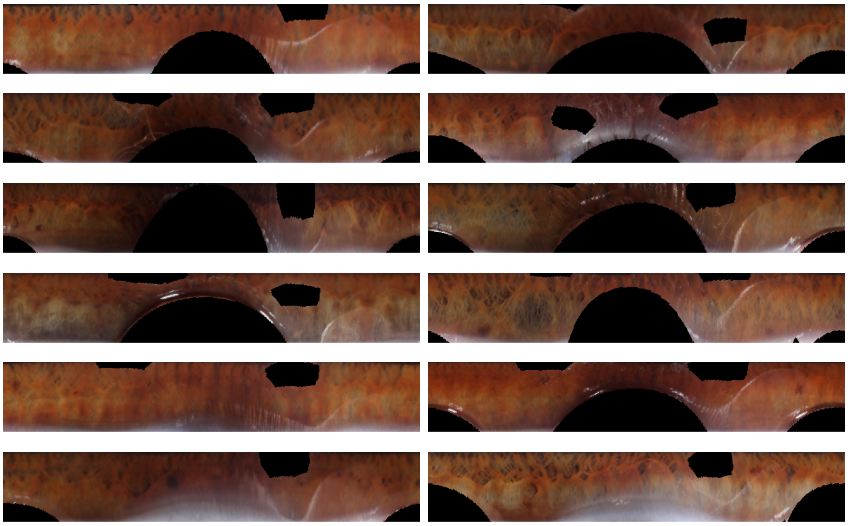

Figure 7.20: Validation images clustered based on colour - group 4

In Figure 7.11 it is seen that the dendrogram for the combination of structure and colour has four large clusters. The number of clusters was chosen based on the appearance of the dendrogram. The iris maps for the groups are seen in Figure 7.21 to 7.26. It is seen that the group 1 contains primarily blue eye images, along with blue eye images containing a small amount of brown. The group seems to contain both eye images with much structure as well as more smooth looking eye images. Group 2 contains intermediate eye colours with much structure. Group 3 also contains intermediate eye colours but have a more smooth appearance. Group 4 contains brown eye images with a small amount of structure present.

The general clustering of the images based on combined structure and colour is very similar to the clustering based on colour only. It is observed that 14 images have changed group in comparison with the colour based clustering. 12 of these 14 images change from group 2 to group 1, which means that they change from light intermediate to blue. The two remaining images change from group 3 to group 2 and from group 4 to group 3, respectively.

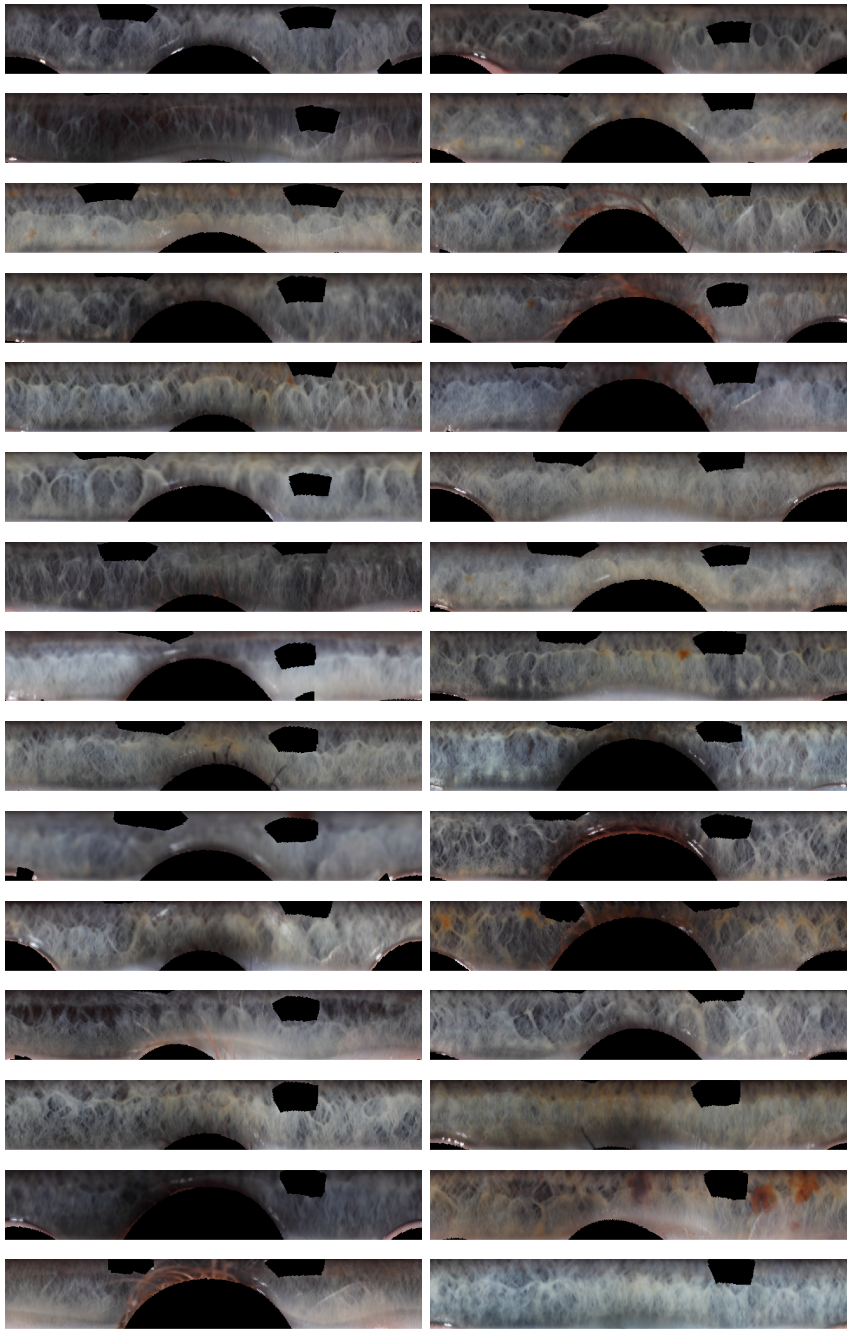

Figure 7.21: Validation images clustered based on a combination of structure and colour - group 1

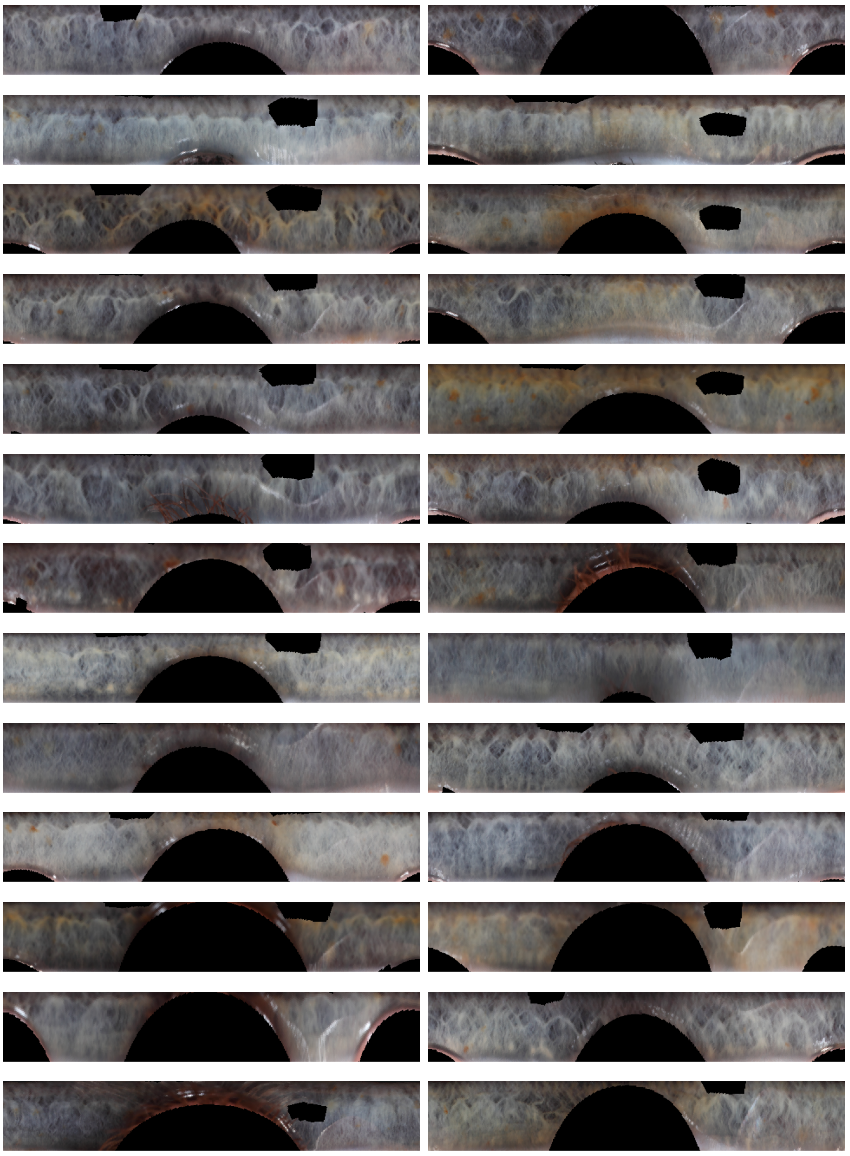

Figure 7.22: Validation images clustered based on a combination of structure and colour - group 1

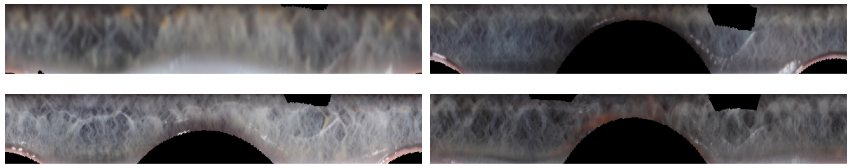

Figure 7.23: Validation images clustered based on a combination of structure and colour - group 1

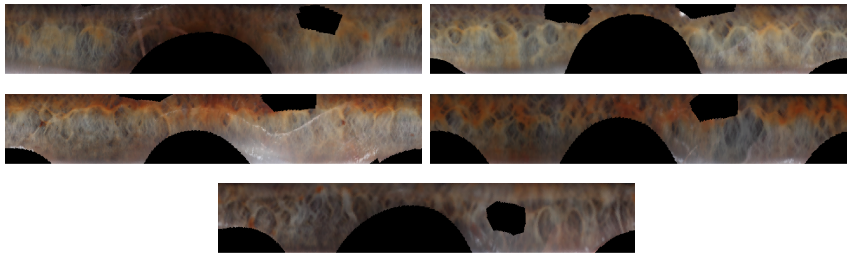

Figure 7.24: Validation images clustered based on a combination of structure and colour - group 2

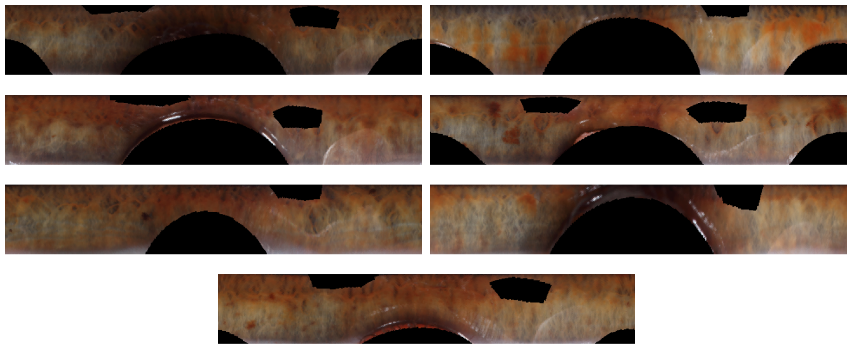

Figure 7.25: Validation images clustered based on a combination of structure and colour - group 3

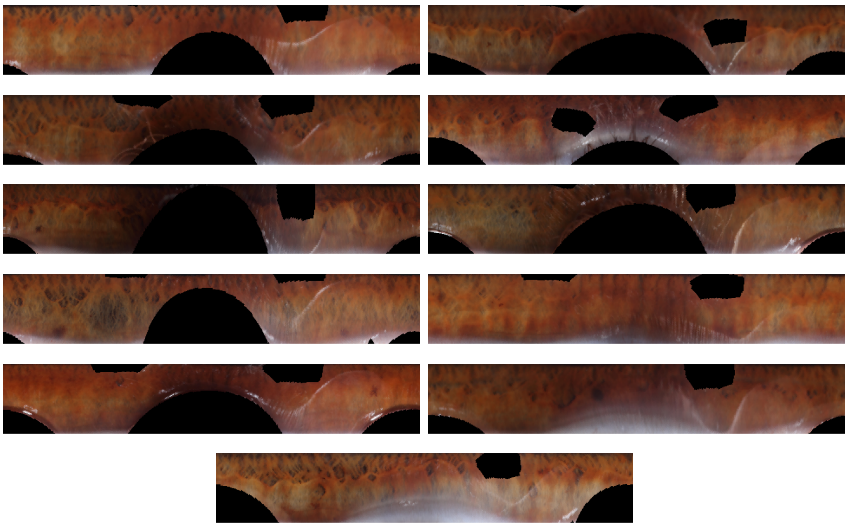

Figure 7.26: Validation images clustered based on a combination of structure and colour - group 4

7.3.1 Colour approach and blue vs brown ratio

The ratios found in Chapter 5 have been divided into the four groups found by the colour based data driven clustering. The boxplot is seen in Figure 7.27. The general behaviour is, that the blue group has a ratio of approximately 1. The light and dark intermediate groups span the interval between blue and brown, but are clearly different with respect to mean ratio. The brown group has a mean ratio of approximately -0.8.

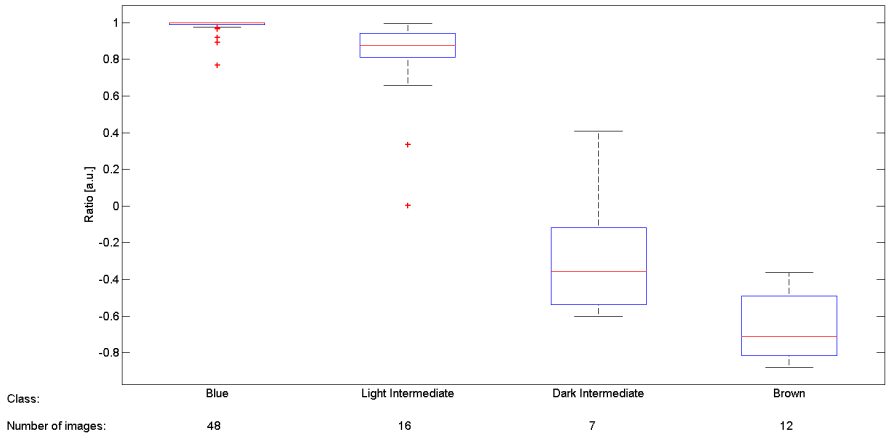

Figure 7.27: Boxplot for the colour based data driven approach

7.4 Image Clustering Discussion

General parameters There are two main parameters that can be adjusted in order to optimize the performance of the clustering method. The first parameter is the number of descriptors used in the training process. The descriptors are calculated at randomly chosen pixel positions in the training image. The performance was tested using 20000, 40000 and 60000 descriptors. The clustering process was not very robust using 20000 descriptors, since the clustering result changed for each new training. The number of descriptors was raised to 40000, which resulted in a more robust training process and thereby more robust clustering. The result for using 60000 descriptors was equivalent to the result with 40000 descriptors, and therefore 40000 descriptors was chosen.

The second parameter is the number of clusters used to build the dictionary of visual words. The algorithm was tested using 200, 400 and 600 clusters. The performance of the algorithm was improved when using 400 clusters instead of 200 clusters. The result for 600 clusters was equivalent to the result for 400 clusters, and 400 clusters were therefore chosen.

Clustering based on structure The size of the DAISY descriptor, illustrated in Figure 7.2, is defined by determining the radius from the center of the descriptor to the center of each circle. The size of the DAISY descriptor has been tested using the set of radii 2, 6, 11 and 6, 12, 25. The clustering result was more intuitive using the largest DAISY descriptor, since the clustering seems to be divided into groups based on the general amount of structure. Therefore the radii were chosen as 6, 12 and 25.

The result of the structure based clustering procedure, is seen in Figure 7.9. The number of cluster groups chosen was three, since more than three groups would result in groups containing very few images. Group 1 contains eye images with a high amount of structure and group 3 contains smooth eye images. On the contrary group 2 seems to be based on a reflection in the eye images and must therefore be considered noise with respect to a DNA correlation.

The division of images into groups containing irises with a high amount of structure and smooth looking irises, respectively, is very promising with respect to DNA correlation. However, there is some images that should be placed differently based on a visual inspection of the cluster groups. The eight images are seen in Figure 7.28 and Figure 7.29 a) to d). They all appear very smooth compared to the remaining in group 1. However, some of the images contains areas that can explain the clustering result. Image 5, 64 and 109, seen in Figure 7.28, all contain eyelash reflections, which are probably interpreted as structure

by the algorithm. Image 22 has a horizontal structure, even though the general appearance is very smooth, and this structure could explain why the image was placed in group 1. Why the clustering procedure has placed the remaining of the eight images into group 1 is not obvious.

The images in Figure 7.29 e) and f) were placed in group 3 by the algorithm. However, the two images contain a high amount of structure compared to the remaining images in group 3. The reason for this is in our opinion misplacement not obvious.

(a) Iris map for image 5

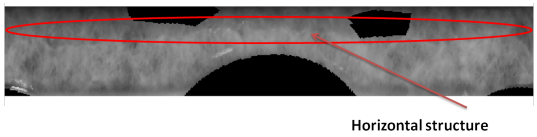

(b) Iris map for image 22

(c) Iris map for image 64

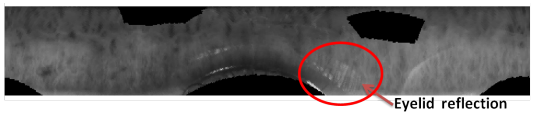

(d) Iris map for image 109

Figure 7.28: Images with explainable misplacement according to a visual inspection. The red circles surrounds the misplacement area.

Clustering based on colour The colour based clustering resulted in a very nice division of the eye images into four groups. The four groups are blue, light intermediate, dark intermediate and brown. The division of the images is much more robust than the colour classifier explained in Chapter 4. Furthermore, this approach is completely data driven and not dependent on subjective evaluations or on previously defined classes.

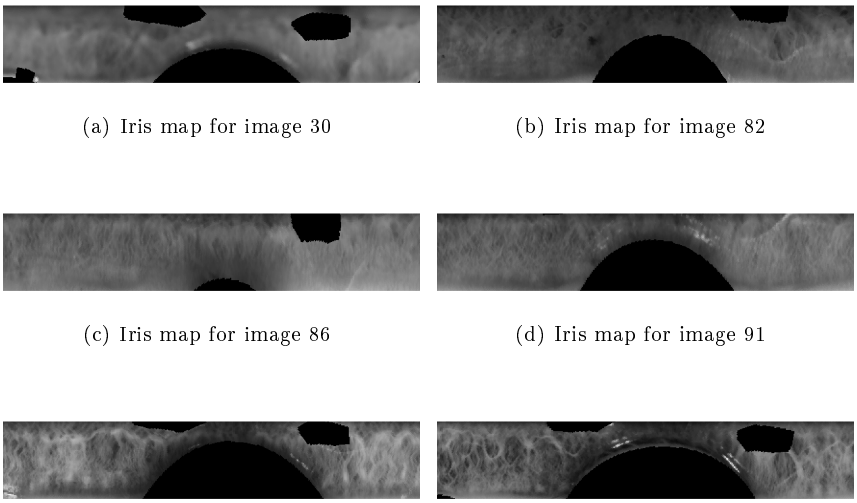

(a) Iris map for image 30

(b) Iris map for image 82

(c) Iris map for image 86

(d) Iris map for image 91

(e) Iris map for image 28

(f) Iris map for image 32

Figure 7.29: Images that are misplaced according to a visual inspection. a) through d) was misplaced in group 1 and e) to f) was misplaced in group 3

The ratios found in Chapter 5 are seen in the boxplot in Figure 7.27, where they have been divided according to the four groups found by the colour based clustering. The boxplot reveals four very clearly differentiable groups. Furthermore, a smooth transition from blue to brown is present, as would be expected, since eye colour is continuous. Comparing this boxplot to the boxplot in Figure 5.4, where the ratios were divided based on subjective evaluations, it is seen that the data driven division of the groups seems much more appropriate. The boxplot using the subjective evaluations has one group that spans the entire interval between brown and blue, and it has multiple classes that are very similar in mean value. The boxplot using the data driven approach divides the images so that no groups overlap.

Clustering based on structure and colour combined The clustering performed on colour and structure combined is very similar to the clustering based on colour. However, there are 14 images that have changed group in comparison with the colour based clustering. 12 of these 14 images changed from group 2 to group 1, meaning that they change from light intermediate to blue. The two

remaining images changed from group 3 to group 2 and from group 4 to group 3, respectively. The conclusion is therefore, that structure does have some influence on the resulting clustering, but the colour restricts the images from moving to a group containing very different colours.

The most promising result obtained using the data driven clustering is the colour based clustering. The clustering result is very robust, and the groups are perfect for correlation with DNA. The structure based clustering result is however more inconclusive, since the results are very difficult to evaluate visually. The clustering into a group containing structure and a group containing more smooth images is very promising. However there are some images that seem to be incorrectly placed. Due to the difficulties with making a subjective evaluation of the structures in the image, it is very difficult to evaluate the result further. We are not convinced that the two groups of structure can give a meaningful correlation with DNA.

The clustering based on a combination of colour and structure also appears to be very promising. The effect from colour versus structure seems to be nicely balanced, since both structure and colour influence the result.

The structure in the eye images could be dependent on the iris colour. If the iris structure was colour dependent, a hypothesis could be that brown irises contained less structure, thereby appearing more smooth. However, this does not seem to be the case for our eye images, since the tree structure based groups contain both blue and brown images. Our validation dataset does not contain any very dark eye images, and therefore it is impossible for us to conclude on the colour dependency of structures. It might be the case that dark eye images contain less structures, and that all of our validation images are simply relatively light eye images.

Chapter 8

Conclusion

The main goal of this thesis was to develop a method for obtaining an objective evaluation of the colours and structures of the iris. Several methods were developed for this purpose. In order to fulfil the goal the iris region had to be extracted from the eye images. We succeeded in extracting the iris region in 64.3% of the eye images of the validation data.

A colour classifier was build using a subjective evaluation of iris colour in the training process. The obtained percentage of correctly classified images was 65.4% for the validation images. Since the colour classifier was based on a histogram with a high amount of dimensions a dimensionality reduction was performed. The six first principal components explained more than 95% of the variance and this number of dimensions was more appropriate for correlation with DNA. However, the six principal components are difficult to interpret.We therefore developed a method to describe a more explanatory ratio. The ratio explained the relationship between the amount of blue and brown in the eye. The behaviour of the ratio corresponds to the behaviour expected from the knowledge about DNA coding for eye colour, since a smooth transition from blue to brown is present. The intermediate eye colours are still to be identified in the genetics, and our ratio is a promising measure for the this purpose, since it clearly distinguish the intermediate groups.

As part of this thesis an investigation of the human perception of eye colour

was performed. The conclusion was that the individuals had a very low degree of agreement for all the investigated iris structures. A classifier could therefore not be trained on subjective evaluations of iris structure.

A data driven image clustering approach was developed in order to account for the problems with human perception of both colour and structure. The result for the structure based clustering was a division into three groups: one containing irises with high amount of structure, a second containing smooth irises and a third containing images with large reflection areas. This division is very promising regarding a DNA correlation. However there were a number of images that were incorrectly placed based on a visual inspection, which makes us question the results. The result from the colour based clustering was very impressing. The four groups created by the method, were clearly divided based on the colour information contained in the iris, since a blue, brown and two intermediate groups were found. This division is more promising since it is independent of subjective evaluations. Furthermore, these groups are more appropriate for explaining the dataset, since no dark brown group is present due to the lack of dark brown images. The result from the combined clustering showed that structure does influence the clustering, but that the colour information restricts the structure dependency. Whether structure and colour is dependent on each other was not possible to conclude based on our given dataset.

In summary, we have succeed in generating a colour classifier, a colour ratio and an image clustering method based on both structure and colour. These methods can be used for making a more objective evaluation of iris colours than what have been published in the literature so far. Furthermore, the results from this thesis might contribute to a further knowledge about the genetics of iris structures and colours.

Chapter 9

Future Work

This thesis concentrates on iris extraction and an evaluation of the colour and structures of the iris region.

The iris extraction process relies on the result from four different tasks, and furthermore it is dependent on the capturing conditions. The iris extraction could be made more robust if one method could replace the four tasks performed in this thesis. A method using textures instead of colour could make the iris extraction less dependent on the image colour and thereby the capturing conditions. An obvious choice would be a texture segmentation method as the one used in [16], commonly used when classifying different structures in for example tissue images.

A very interesting field would be to make a colour classifier trained on objectively classified groups. A future work could be to combine the results from the data driven clustering in Chapter 7 with the colour classifier in Chapter 4, in order to create a completely objective colour classifier. The approach would be to use the four colour classes found by the data driven clustering as training set in the colour classifier.

A pixel classification of the iris region is of interest in order to locate the different structures in the iris. The structures of interest are for example Nevi Dots, Wolfflin nodules and Fuchs Crypts. Nevi Dots are of great interest regarding

risk of developing cancer, and a method for extracting Nevi Dots is therefore very desired. The pixel classification using MRF, as in Chapter 5 to identify brown and blue areas, might be a potential method for finding Nevi Dots in blue eyes. The extraction of brown areas could be the first step, followed by an evaluation of the structures in the extracted brown spots. The Nevi Dots have a more smooth appearance and does not follow the general structures in the iris region, why a texture classification method could be appropriate.

Appendix A

Iris Extraction

A.1 Optimization of Processing Time

The most obvious method to use when searching for a local minimum in a gradient function is the method of Steepest Descent, which is a fast method to locate a local minimum. Potentially, this method might decrease the computational effort and time aspect in our implementation. The Steepest Descent method is an acknowledged method and it has been described by many authors, e.g. Poul Erik Frandsen et. al. [10]. The method searches for a local minimum of a given function. The search direction is determined as the direction in which the gradient has the highest negative value. Searching for a minimum is the same approach as searching for a maximum, just with a sign change. In the following we will search for the maximum.

In order to test the implementation, a simple circle image is generated, see Figure A.1 a). In order to use the Steepest Descent method the function to maximize should be relatively smooth. The function to maximize is, in our case, a function with three variables, center coordinate (x,y) and radius, r. Figure A.1 b) shows a plot of this function. In the figure a maximum gradient sum is found for each pixel. For all pixels the gradient sum is calculated for all radii, and the radius with the maximum gradient is chosen. It is seen that the function has many local maxima. These local maxima makes the task of locating the global maximum

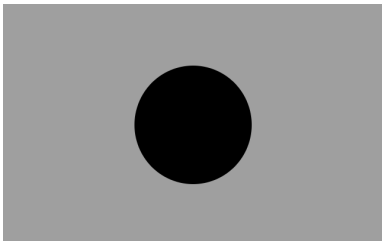

(a) Original circle image

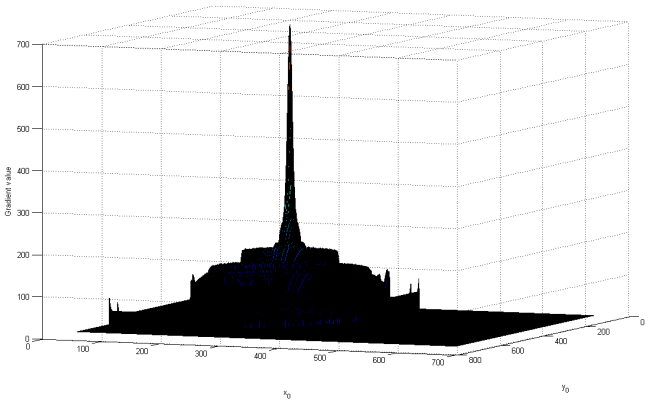

(b) Gradient values

Figure A.1: a) Circle image b) Function to minimize. The gradient values on the longitudinal axis are the maximum of the gradient sums for all radii for that given center coordinate.

more difficult. Using the steepest descent method to locate the global maximum would create problems, since the method would stop searching as soon as a local maximum was located.

A.2 Graph Cut

Given the system, illustrated in Figure A.2 top left, the first step is to find the maximum flow through the system, and then transfer the minimal capacity through the system. The path of maximum flow is lowered by the flow size. As seen in Figure A.2 a) top right, the path $s-1-3-t$ has been saturated, meaning that nothing can be transported through this specific path. Continuing this

procedure, until it is not possible to transfer energy trough the system, results in a saturated system, and the cut is placed at the position of saturation, see Figure A.2 a) bottom right. As is seen from the figure, the max flow, min cut procedure, separates the nodes 1, 2 and 4 from 3.

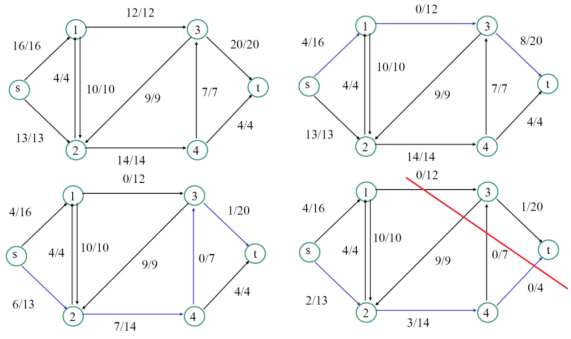

Figure A.2: The figure illustrates the Graph Cut method [9]

A.3 Iris Extraction Results

A sample of the results from the iris extraction method are seen in Figure A.3 and A.4. The images are defined as correctly detected, since they have a Dice Coefficient above 0.92, as chosen in Section 3.4.

Some examples of the iris extraction result are seen in Figure A.5. They are defined as wrongly detected, since the Dice Coefficient is below 0.92. It is seen that some of the wrongly detected images actually have satisfying detected boundaries. They could therefore be included in the further methods in our thesis, as stated in Section 3.7.

[9] Lecture slides from course Advanced Image Analysis 02503, spring 2011

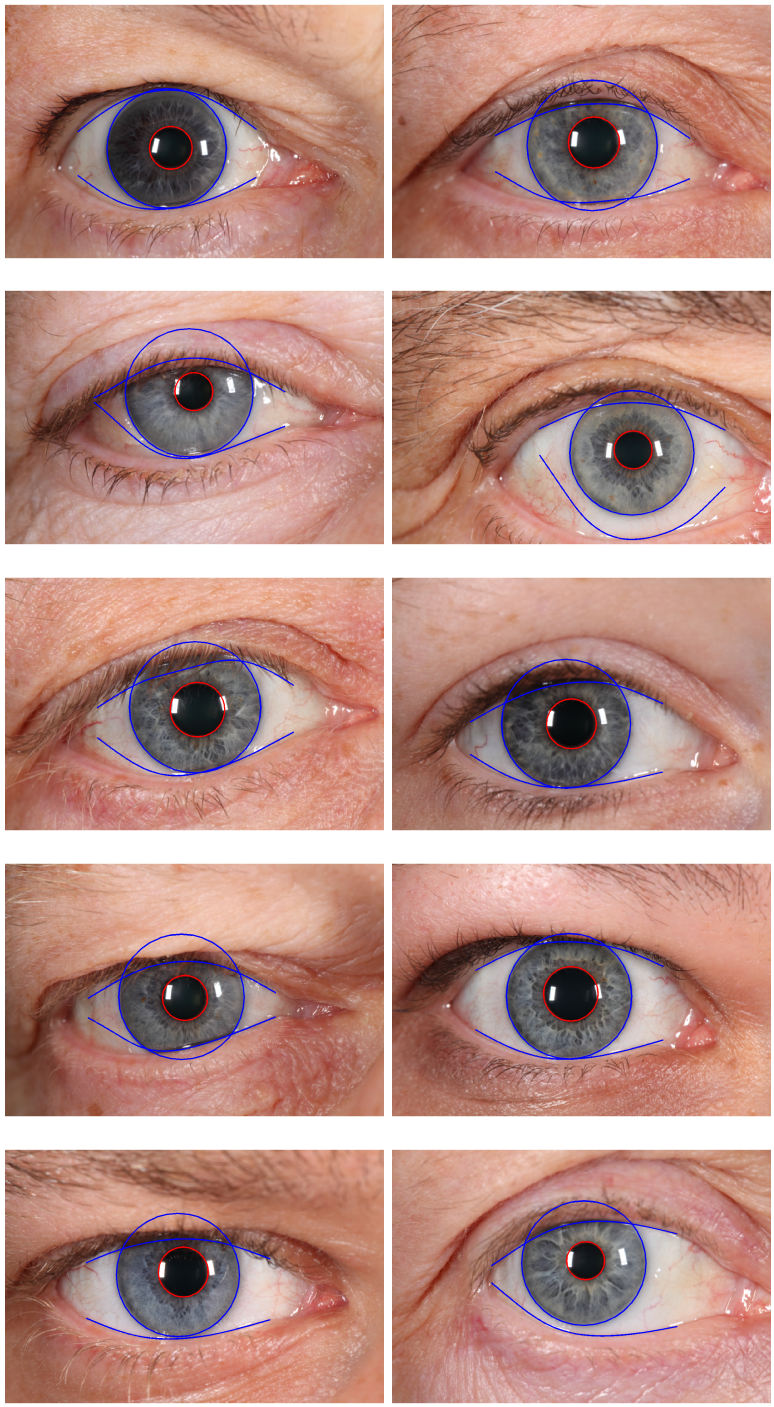

Figure A.3: Eye images with correctly detected iris regions

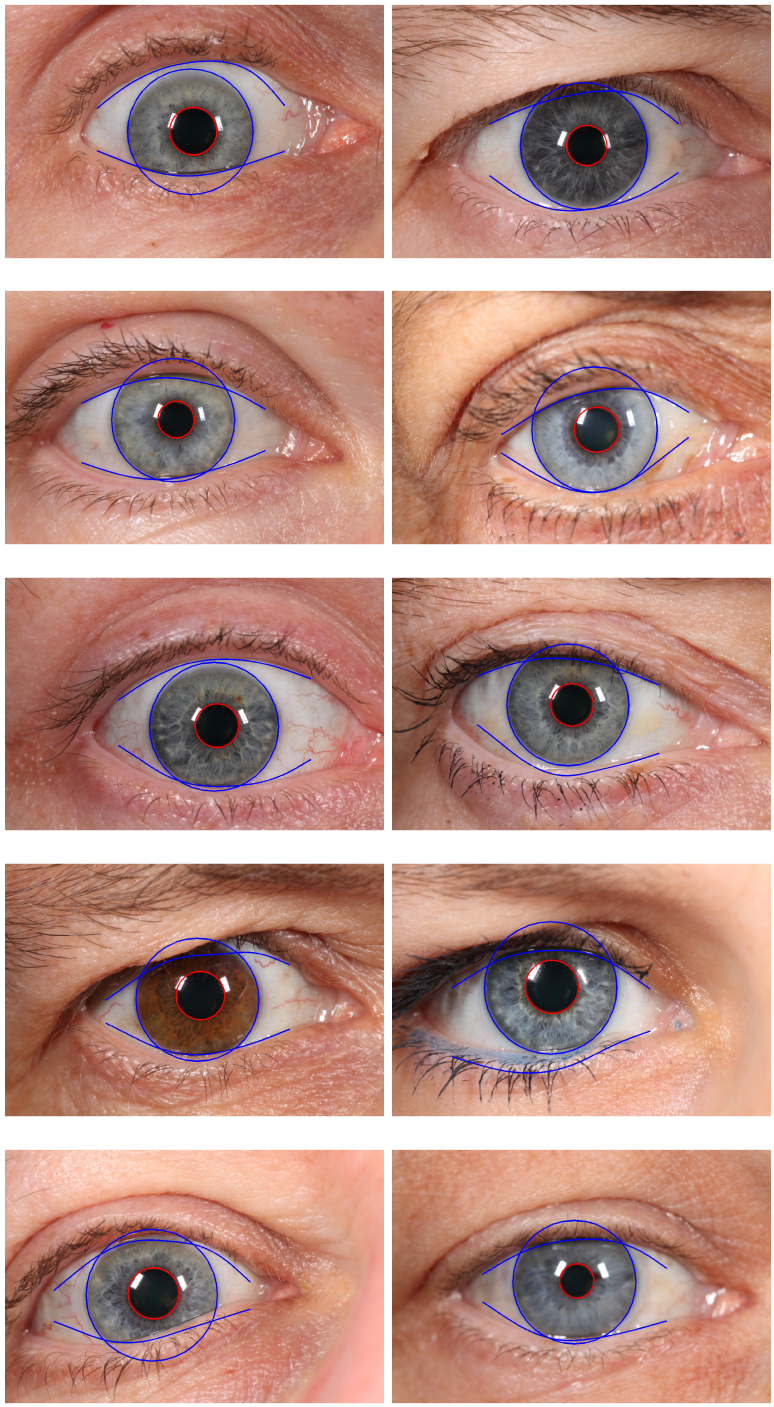

Figure A.4: Eye images with correctly detected iris regions

Figure A.5: Eye images with wrongly detected iris regions

A.4 Iris Map Generation Results

A small selection of iris maps, generated from the radial transformation, and the corresponding iris masks are seen in Figure A.6. They are all defined as correctly detected from the iris extraction.

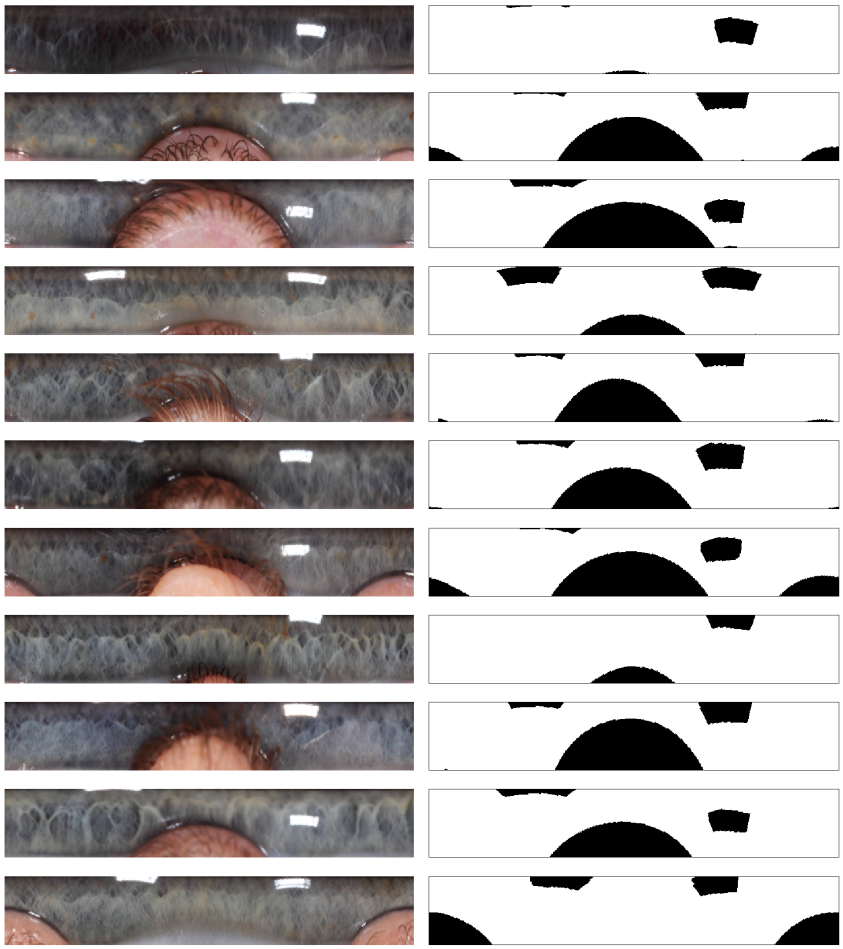

Figure A.6: Iris maps and corresponding iris masks, for validation images with correctly detected iris regions

Appendix B

Colour Classifier

B.1 Human Evaluation of Eye Colour Classes

The subjective evaluations of the images have been performed as described in Section 4.1. The evaluations of the 100 training images are seen in Table B.1 to B.3, where yes and no indicate if an agreement or a disagreement is present. The percentage of disagreement between the two operators is 32%.

Image number	Individual 1	Individual 2	Comparison
1	2	2	yes
2	1	1	yes
3	1	1	yes
4	1	1	yes
5	2	2	yes
6	4	4	yes
7	5	6	no
8	2	2	yes
9	2	6	no
10	6	6	yes

Table B.1: Human evaluation of colour classes for the training images

Image number	Individual 1	Individual 2	Comparison
11	4	4	yes
12	4	6	no
13	4	4	yes
14	2	2	yes
15	2	2	yes
16	1	1	yes
17	3	4	no
18	2	2	yes
19	6	2	no
20	3	3	yes
21	6	6	yes
22	2	2	yes
23	5	5	yes
24	2	1	no
25	4	4	yes
26	5	5	yes
27	2	2	yes
28	2	2	yes
29	5	1	no
30	5	1	no
31	1	1	yes
32	6	2	no
33	6	6	yes
34	5	1	no
35	2	2	yes
36	2	2	yes
37	2	2	yes
38	6	5	no
39	2	2	yes
40	2	2	yes
41	2	6	no
42	2	2	yes
43	6	5	no
44	5	1	no
45	6	2	no
46	5	5	yes
47	2	2	yes
48	2	1	no
49	2	1	no
50	4	4	yes
51	1	1	yes
52	5	5	yes
53	6	5	no
54	5	1	no
55	6	6	yes
56	2	2	yes
57	4	6	no
58	1	1	yes
59	2	2	yes
60	2	2	yes

Table B.2: Human evaluation of colour classes for the training images

Image number	Individual 1	Individual 2	Comparison
61	2	2	yes
62	1	1	yes
63	2	1	no
64	6	2	no
65	2	2	yes
66	3	3	yes
67	4	3	no
68	2	2	yes
69	1	1	yes
70	5	5	yes
71	6	6	yes
72	6	6	yes
73	2	2	yes
74	2	2	yes
75	2	1	no
76	1	1	yes
77	2	1	no
78	2	1	no
79	3	3	yes
80	1	1	yes
81	1	1	yes
82	2	2	yes
83	3	6	no
84	1	1	yes
85	1	1	yes
86	2	2	yes
87	2	2	yes
88	3	3	yes
89	6	2	no
90	5	5	yes
91	1	1	yes
92	6	3	no
93	2	2	yes
94	2	1	no
95	1	1	yes
96	4	3	no
97	2	2	yes
98	1	1	yes
99	5	5	yes
100	6	3	no

Table B.3: Human evaluation of colour classes for the training images

B.2 Choosing a Colour Classifier

In order to determine which classifier method to use for the final data run, we have tested the different methods for k = 1...5, where k is the number of k-nearest images. If there are multiple classes with the same number of images

among the k-nearest images, the class is chosen randomly among the possible classes. In order to account for the randomness, the methods are tested in 10 runs. The results from the tests can be seen in Table B.4 through B.6, for RGB space, LAB space and the histogram approach in RGB space, respectively. An average for each k = 1...5 has been calculated in order to decide upon a method.

	k	1	2	3	4	5
Run 1	Error	25	24	21	19	19
	Correct	23	24	27	29	29
	Correct (%)	47.9	50	56.3	60.4	60.4
Run 2	Error	-	26	22	21	20
	Correct	-	22	26	27	28
	Correct (%)	-	45.8	54.2	56.3	58.3
Run 3	Error	-	27	22	20	19
	Correct	-	21	26	28	29
	Correct (%)	-	43.8	54.2	58.3	60.4
Run 4	Error	-	24	22	21	19
	Correct	-	24	26	27	29
	Correct (%)	-	50	54.2	56.3	60.4
Run 5	Error	-	22	21	22	20
	Correct	-	26	27	26	28
	Correct (%)	-	54.2	56.3	54.2	58.3
Run 6	Error	-	21	22	21	19
	Correct	-	27	26	27	29
	Correct (%)	-	56.3	54.2	56.3	60.4
Run 7	Error	-	22	21	20	19
	Correct	-	26	27	28	29
	Correct (%)	-	54.2	56.3	58.3	60.4
Run 8	Error	-	21	21	21	19
	Correct	-	27	27	27	29
	Correct (%)	-	56.3	56.3	56.3	60.4
Run 9	Error	-	23	22	22	19
	Correct	-	25	26	26	29
	Correct (%)	-	52.1	54.2	54.2	60.4
Run 10	Error	-	23	22	20	20
	Correct	-	25	26	28	28
	Correct (%)	-	52.1	54.2	58.3	58.3
Average	Error	25.0	23.3	21.6	20.7	19.3
	Correct	23.0	24.7	26.4	27.3	28.7
	Correct (%)	47.9	51.5	55.0	56.9	59.8

Table B.4: Testing the RGB space method on 10 runs to obtain an average performance for each k

From the tests it is seen that the highest average percentage is given by the histogram approach for k = 2. We are therefore using this method with the respective k on the 83 validation images. In order to get an average performance of the 83 validation images, we have run the method 10 times. The result can be seen in Table B.7, where it is seen that the final percentage of correctly classified images is 65.4%.

	k	1	2	3	4	5
Run 1	Error	29	29	18	16	21
	Correct	19	19	30	32	27
	Correct (%)	39.6	39.6	62.5	66.7	56.3
Run 2	Error	-	26	18	18	21
	Correct	-	22	30	30	27
	Correct (%)	-	45.8	62.5	62.5	56.3
Run 3	Error	-	30	16	18	21
	Correct	-	18	32	30	27
	Correct (%)	-	37.5	66.7	62.5	56.3
Run 4	Error	-	24	18	18	20
	Correct	-	24	30	30	28
	Correct (%)	-	50	62.5	62.5	58.3
Run 5	Error	-	26	17	16	20
	Correct	-	22	31	32	28
	Correct (%)	-	45.8	64.6	66.7	58.3
Run 6	Error	-	29	17	18	20
	Correct	-	19	31	30	28
	Correct (%)	-	39.6	64.6	62.5	58.3
Run 7	Error	-	29	18	18	19
	Correct	-	19	30	30	29
	Correct (%)	-	39.6	62.5	62.5	60.4
Run 8	Error	-	28	18	18	19
	Correct	-	20	30	30	29
	Correct (%)	-	41.7	62.5	62.5	60.4
Run 9	Error	-	31	17	20	21
	Correct	-	17	31	28	27
	Correct (%)	-	35.4	64.6	58.3	56.3
Run 10	Error	-	22	18	19	20
	Correct	-	26	30	29	28
	Correct (%)	-	54.2	62.5	60.4	58.3
Average	Error	29	27.4	17.5	17.9	20.2
	Correct	19	20.6	30.5	30.1	27.8
	Correct (%)	39.6	42.9	63.5	62.7	57.9

Table B.5: Testing the LAB space method on 10 runs to obtain an average performance for each k

	k	1	2	3	4	5
Run 1	Error	13	4	12	12	16
	Correct	35	44	36	36	32
	Correct (%)	72.9	91.7	75.0	75.0	66.7
Run 2	Error	-	7	13	15	15
	Correct	-	41	35	33	33
	Correct (%)	-	85.4	72.9	68.8	68.8
Run 3	Error	-	12	13	17	12
	Correct	-	36	35	31	36
	Correct (%)	-	75	72.9	64.6	75
Run 4	Error	-	12	12	14	14
	Correct	-	36	36	34	34
	Correct (%)	-	75.0	75.0	70.8	70.8
Run 5	Error	-	14	15	10	12
	Correct	-	34	36	38	36
	Correct (%)	-	70.8	75.0	79.2	75
Run 6	Error	-	11	12	13	16
	Correct	-	37	36	35	32
	Correct (%)	-	77.1	75.0	72.9	66.7
Run 7	Error	-	13	12	8	15
	Correct	-	35	36	40	33
	Correct (%)	-	72.9	75.0	83.3	68.8
Run 8	Error	-	14	12	10	12
	Correct	-	34	36	38	36
	Correct (%)	-	70.8	75.0	79.2	75.0
Run 9	Error	-	7	12	13	16
	Correct	-	41	36	35	32
	Correct (%)	-	85.4	75.0	72.9	66.7
Run 10	Error	-	13	12	15	14
	Correct	-	35	36	33	34
	Correct (%)	-	72.9	75.0	68.8	70.8
Average	Error	13	10.7	12.2	12.7	14.2
	Correct	35	37.3	35.8	35.3	33.8
	Correct (%)	72.9	77.7	74.6	73.5	70.4

Table B.6: Testing the Final method based on RGB space histograms on 10 runs to obtain an average performance for each k

	k	2
Run 1	Error	30
	Correct	53
	Correct (%)	63.9
Run 2	Error	29
	Correct	54
	Correct (%)	65.1
Run 3	Error	33
	Correct	50
	Correct (%)	60.2
Run 4	Error	31
	Correct	52
	Correct (%)	62.7
Run 5	Error	32
	Correct	51
	Correct (%)	61.4
Run 6	Error	24
	Correct	59
	Correct (%)	71.1
Run 7	Error	28
	Correct	55
	Correct (%)	66.3
Run 8	Error	29
	Correct	54
	Correct (%)	65.1
Run 9	Error	26
	Correct	57
	Correct (%)	68.7
Run 10	Error	25
	Correct	58
	Correct (%)	69.8
Average	Error	28.7
	Correct	54.3
	Correct (%)	65.4

Table B.7: Testing the Final method based on RGB space histograms on 10 runs to obtain an average performance for each k

In order to get an overview of where the algorithm fails, the subjective evaluations are given in Table B.8 through B.10 along with the classes found for each image by the algorithm. Where the algorithms fails, a no is written, and where it is correct, a yes is written. The results are obtained for a run with 67.5% correctly classified images.

Image number	Subjective evaluation	Algorithm	Comparison
1	2	2	yes
2	2	2	yes
3	6	-	-
4	2	2	yes
5	3	4	no
6	5	5	yes
7	1	-	-
8	-	-	-
9	1	1	yes
10	2	2	yes
11	2	2	yes
12	2	2	yes
13	4	4	yes
14	2	2	yes
15	1	2	no
16	1	1	yes
17	6	-	-
18	6	6	yes
19	2	5	no
20	2	2	yes
21	4	4	yes
22	1	1	yes
23	6	-	-
24	1	2	no
25	2	2	yes
26	2	2	yes
27	4	4	yes
28	2	2	yes
29	1	-	-
30	2	2	yes
31	4	4	yes
32	2	2	yes
33	1	1	yes
34	6	6	yes
35	6	5	no
36	2	2	yes
37	1	-	-
38	4	-	-
39	2	-	-
40	3	6	no
41	6	5	no
42	1	1	yes
43	4	6	no
44	2	-	-
45	4	6	no

Table B.8: Subjective evaluation and algorithm result for the validation images

Image number	Subjective evaluation	Algorithm	Comparison
46	2	-	-
47	3	3	yes
48	1	5	no
49	6	2	no
50	2	2	yes
51	6	4	no
52	5	2	no
53	1	2	no
54	1	1	yes
55	4	6	no
56	2	2	yes
57	1	-	-
58	2	2	yes
59	1	-	-
60	1	-	-
61	1	2	no
62	6	6	yes
63	2	-	-
64	3	4	no
65	4	3	no
66	1	-	-
67	1	1	yes
68	6	-	-
69	2	6	no
70	1	5	no
71	2	2	yes
72	2	-	-
73	2	5	no
74	4	-	-
75	2	2	yes
76	6	6	yes
77	1	-	-
78	2	2	yes
79	2	2	yes
80	2	2	yes
81	2	2	yes
82	3	3	yes
83	1	1	yes
84	2	-	-
85	6	6	yes
86	1	2	no
87	2	-	-
88	4	4	yes
89	1	-	-
90	3	4	no
91	2	2	yes
92	2	-	-
93	2	-	-
94	1	1	yes
95	1	1	yes
96	3	6	no
97	1	2	no
98	2	-	-
99	6	-	-

Table B.9: Subjective evaluation and algorithm result for the validation images

Image number	Subjective evaluation	Algorithm	Comparison
100	2	6	no
101	5	2	no
102	1	1	yes
103	1	1	yes
104	2	2	yes
105	5	5	yes
106	2	2	yes
107	2	2	yes
108	1	1	yes
109	6	3	no
110	1	-	-
111	2	2	yes
112	4	-	-

Table B.10: Subjective evaluation and algorithm result for the validation images

Dimensionality reduction considerations The results from the classifier should be used in correlation with DNA information, as described in Section 4.3. Therefore the classifier has been run on the histograms generated from the six principal components, see Section 4.4. The result from running the classifier on the back transformed histograms is shown in Table B.11. It is seen that the percentage of correctly detected images is 60.2%, which is lower than for the results using the original histograms.

	k	2
Run 1	Error	30
	Correct	53
	Correct (%)	63.9
Run 2	Error	34
	Correct	49
	Correct (%)	59.0
Run 3	Error	35
	Correct	48
	Correct (%)	57.8
Run 4	Error	34
	Correct	49
	Correct (%)	59.0
Run 5	Error	31
	Correct	52
	Correct (%)	62.7
Run 6	Error	36
	Correct	47
	Correct (%)	56.6
Run 7	Error	35
	Correct	48
	Correct (%)	57.8
Run 8	Error	31
	Correct	52
	Correct (%)	62.7
Run 9	Error	33
	Correct	50
	Correct (%)	60.2
Run 10	Error	31
	Correct	52
	Correct (%)	62.7
Average	Error	33
	Correct	50
	Correct (%)	60.2

Table B.11: Testing the Final method based on RGB space histograms with dimensionality reduction on 10 runs to obtain an average performance for each k

Appendix C

Blue vs Brown Ratio

A selection of the results from the pixel classification using MRF is seen in Figure C.1 and C.2. The two classes are brown and blue.

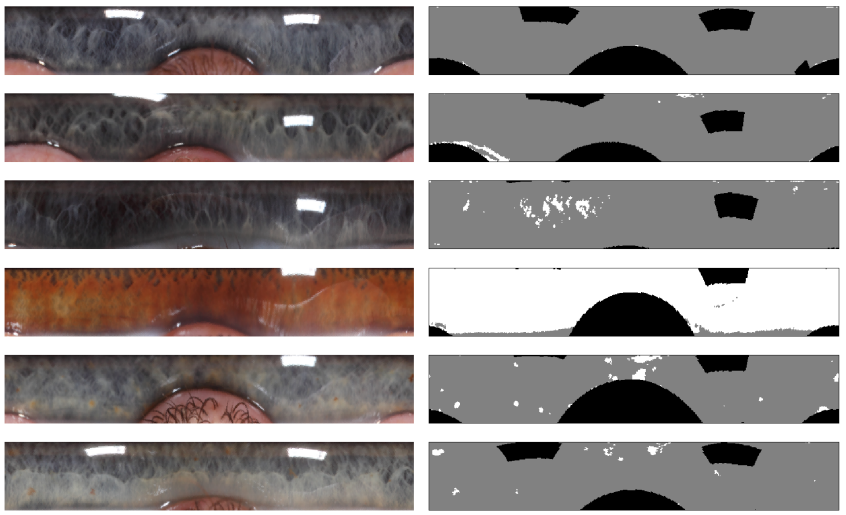

Figure C.1: Results from the brown vs blue pixel classification using MRF

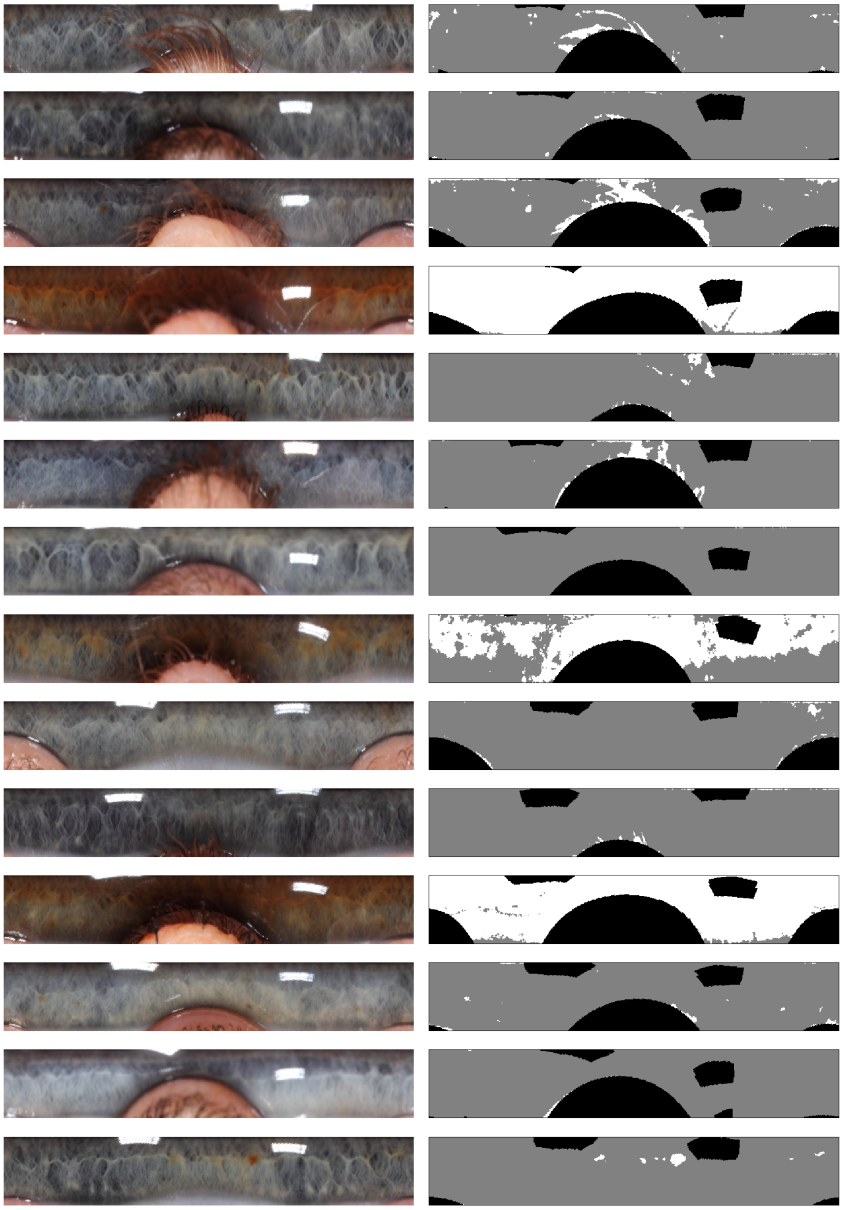

Figure C.2: Results from the brown vs blue pixel classification using MRF

Appendix D

Image Clustering

The result for the structure based image clustering is seen in Figure D.1 to D.5 in colour.

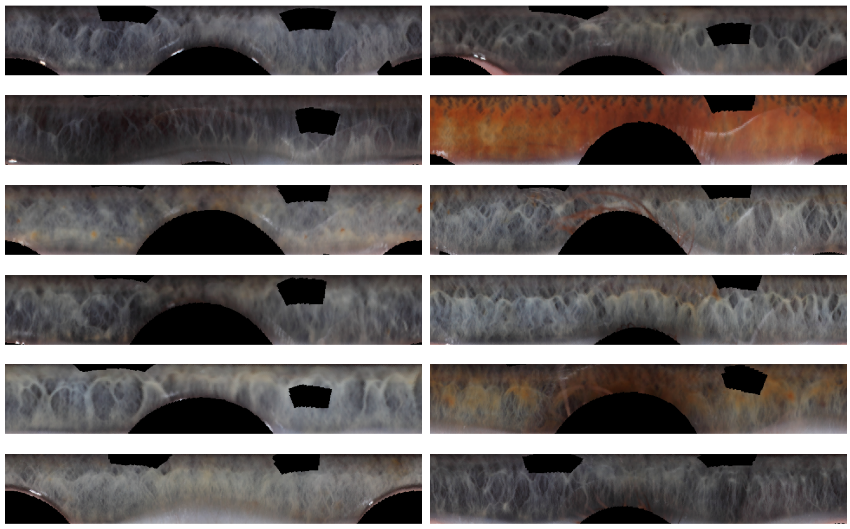

Figure D.1: Validation images clustered based on structure - group 1

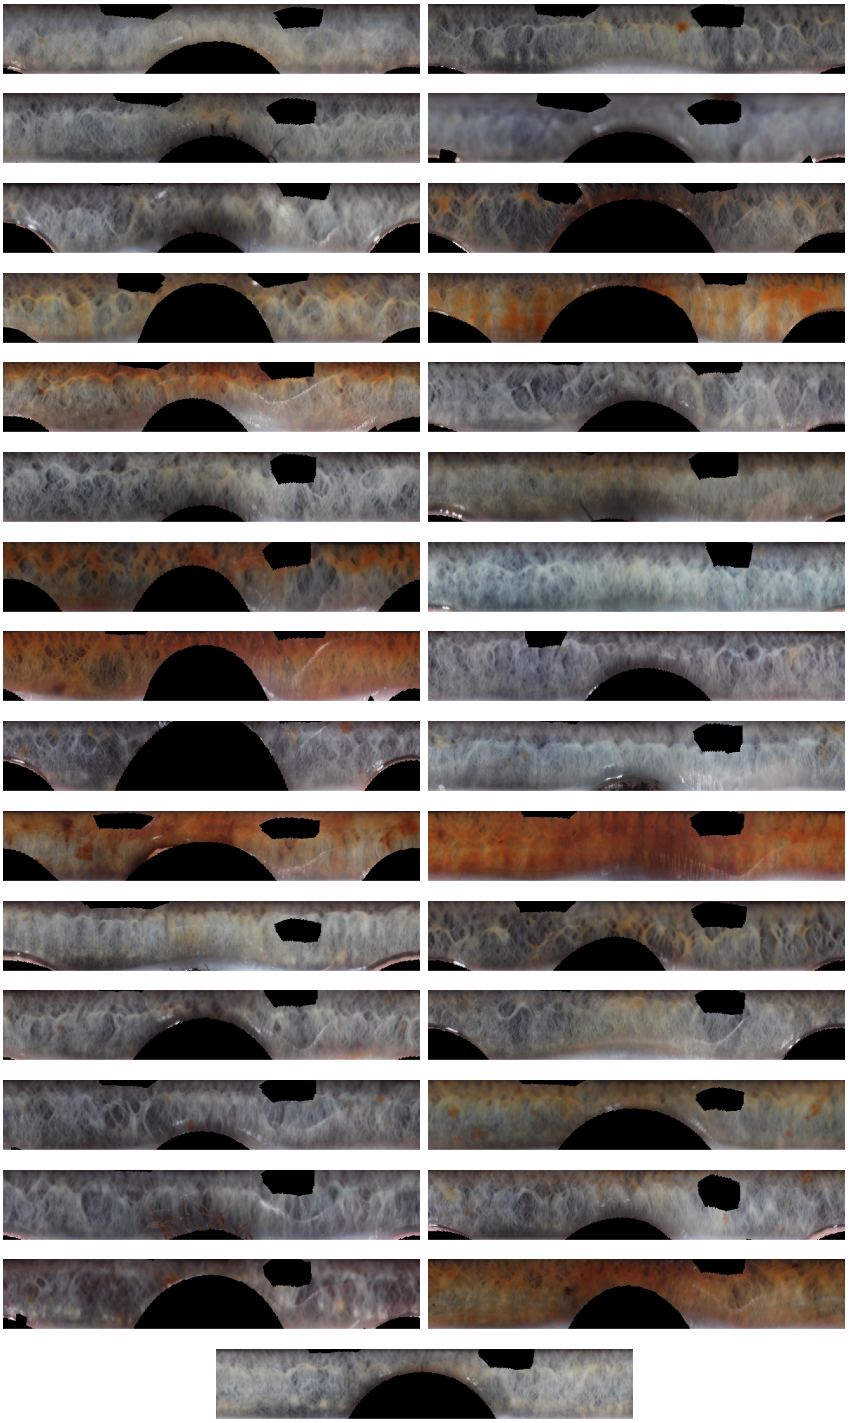

Figure D.2: Validation images clustered based on structure - group 1

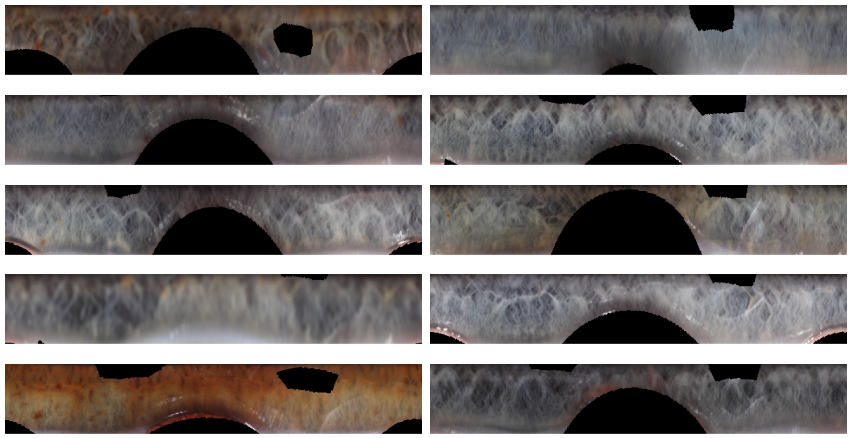

Figure D.3: Validation images clustered based on structure - group 1

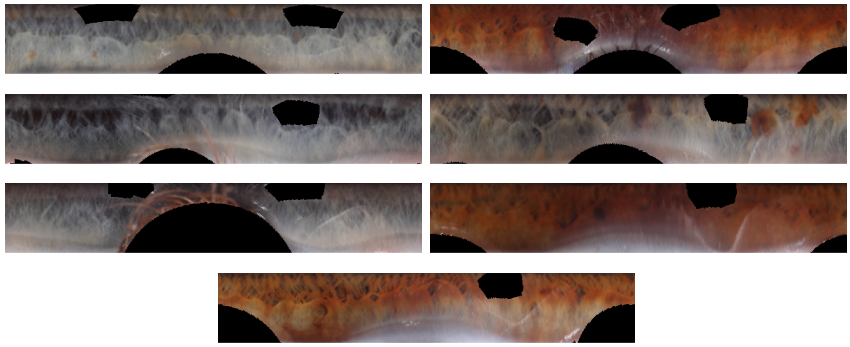

Figure D.4: Validation images clustered based on structure - group 2

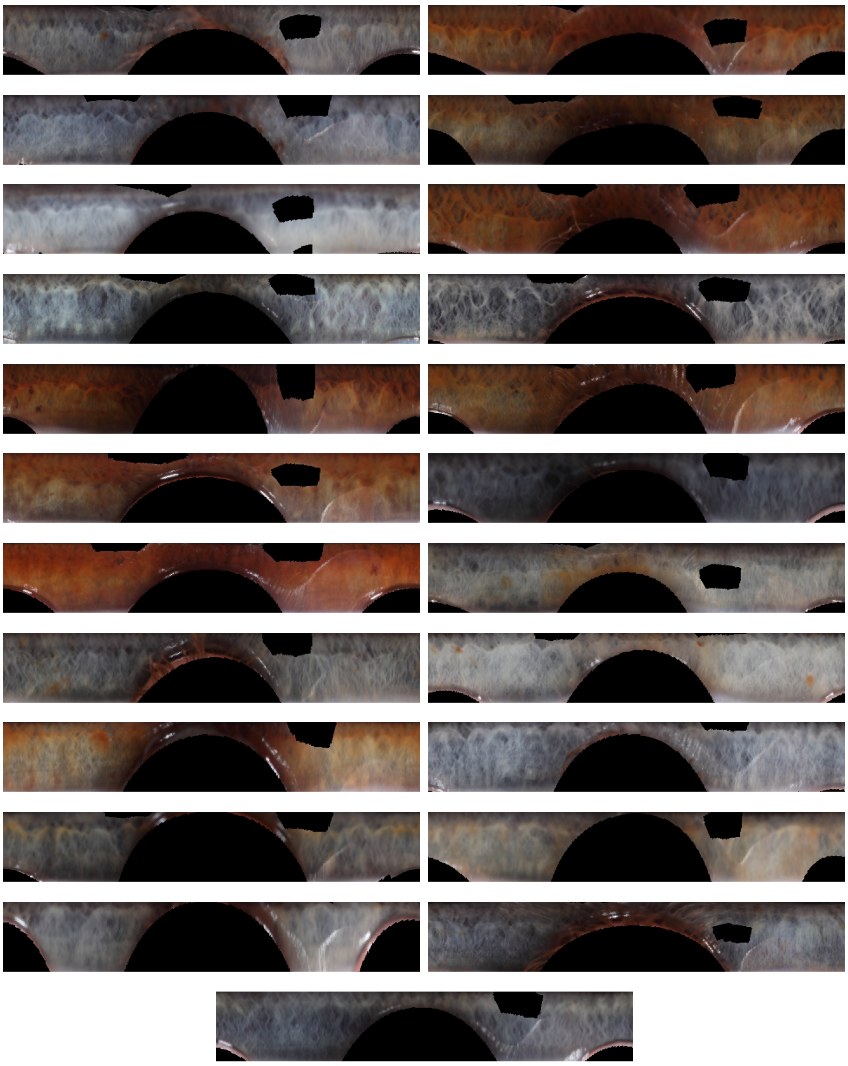

Figure D.5: Validation images clustered based on structure - group 3

Appendix E

CD-ROM Content

In this Appendix a short overview of the CD-ROM content i presented. Furthermore, a copy of the thesis is also included.

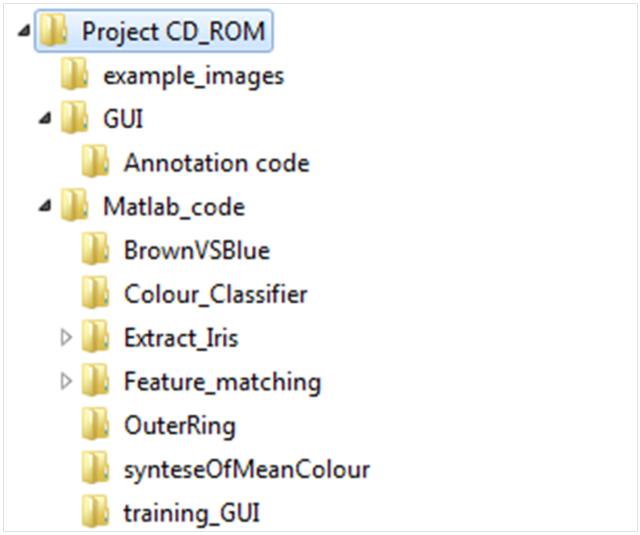

Figure E.1: Contents of the CD-ROM provided with this thesis

Acronyms

Acronym	Explanation
CIELAB	Commision Internationale de L'Eclairage, $L^*a^*b^*$ colour representation
DNA	Deoxyribonucleic acid
DPP	Digital Photo Professional
EVC	External visible characteristics
GUI	Graphical user interface
HSV	Hue, Saturation, Value colour representation
IPE	Iris pigmentation epithelium
MRF	Markov Random Field
PCA	Principal Component Analysis
PNG	Portable Network Graphics
RGB	Red, Green, Blue colour representation
SIFT	Scale Invariant Feature Transform

Table E.1: List of Acronyms

Project Poster

Department of Forensic Medicine,
Section of Forensic Genetics
Danmarks Tekniske Universitet
DTU
CATCH THE BAD GUY
BY SUSANNE R. CHRISTOFFERSEN AND STINE HARDER
INTRODUCTION
Whos the bad guy?
IRIS DETECTION
Method:
Results
COLOUR CLASSIFIER
Method:
Results
SYNTHESIZING
Method:
Results
BROWN VS BLUE DETECTION
Method:
Results

Bibliography

[1] B. Alberts, D. Bray, A. Johnson, J. Lewis, M. Raff, K. Roberts, P. Walter, and A.M. Campbell. *Essential cell biology, 2. th Edition*. Garland Science New York, 2004.

[2] J.L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.

[3] C.M. Bishop and SpringerLink (Service en ligne). *Pattern recognition and machine learning*, volume 4. springer New York, 2006.

[4] M.C. Campbell and S.A. Tishkoff. The evolution of human genetic and phenotypic variation in africa. *Current Biology*, 20(4):R166–R173, 2010.

[5] J. Daugman. How iris recognition works. *Circuits and Systems for Video Technology, IEEE Transactions on*, 14(1):21–30, 2004.

[6] J. Daugman. Probing the uniqueness and randomness of iriscodes: Results from 200 billion iris pair comparisons. *Proceedings of the IEEE*, 94(11):1927–1935, 2006.

[7] J.G. Daugman. High confidence visual recognition of persons by a test of statistical independence. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 15(11):1148–1161, 1993.

[8] H. Deng and D.A. Clausi. Unsupervised image segmentation using a simple mrf model with a new implementation scheme. *Pattern recognition*, 37(12):2323–2335, 2004.

[9] M. Edwards, A. Gozdzik, K. Ross, J. Miles, and E.J. Parra. Technical note: Quantitative measures of iris color using high resolution photographs. *American Journal of Physical Anthropology*, 2011.

[10] Poul Erik Frandsen [et al.]. *Unconstrained optimization, 3th Edition*. Informatics and Mathematical modelling, Technical University of technology, March 2004.

[11] J.L. Fleiss. The measurement of interrater agreement. *Statistical methods for rates and proportions*, 2:212–236, 1981.

[12] T.N. Frudakis. *Molecular photofitting: predicting ancestry and phenotype using DNA*. Academic Press, 2008.

[13] L.K. Hansen, J. Larsen, and T. Kolenda. On independent component analysis for multimedia signals. *Multimedia Image and Video Processing*, pages 175–199, 2000.

[14] R.A. Johnson, I. Miller, and J.E. Freund. *Probability and statistics for engineers*. Prentice-Hall, 2011.

[15] Fan Kuo-Chin Jou, Fan-Di and Yang-Lang Chang. Efficient matching of large-size histograms. *Pattern Recognition Letters*, 25:277–286, 2004.

[16] A. Kårsnäs, A. Dahl, and R. Larsen. Learning histopathological patterns, 2011.

[17] M. Kayser and P. de Knijff. Improving human forensics through advances in genetics, genomics and molecular biology. *Nature Reviews Genetics*, 12(3):179–192, 2011.

[18] P. Kohli and P.H.S. Torr. Dynamic graph cuts for efficient inference in markov random fields. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(12):2079–2088, 2007.

[19] S. Krishnamachari and M. Abdel-Mottaleb. Image browsing using hierarchical clustering. In *Computers and Communications, 1999. Proceedings. IEEE International Symposium on*, pages 301–307. IEEE, 1999.

[20] A. Lefohn, B. Budge, P. Shirley, R. Caruso, and E. Reinhard. An ocularist's approach to human iris synthesis. *Computer Graphics and Applications, IEEE*, 23(6):70–75, 2003.

[21] D.G. Lowe. Object recognition from local scale-invariant features. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 2, pages 1150–1157. Ieee, 1999.

[22] M. Makitalo and A. Foi. On the inversion of the anscombe transformation in low-count poisson image denoising. In *Local and Non-Local Approximation in Image Processing, 2009. LNLA 2009. International Workshop on*, pages 26–32. IEEE, 2009.

[23] J. Mengel-From, C. Børsting, J.J. Sanchez, H. Eiberg, and N. Morling. Human eye colour and herc2, oca2 and matp. *Forensic Science International: Genetics*, 4(5):323–328, 2010.

[24] T.B. Moeslund. *Image and Video Processing*. Computer Vision and Media Technology, Aalborg University, 2008.

[25] R. R Paulsen and B. T Moeslund. *Introduction to Medical Image Analysis*. 2011.

[26] E.E. Quillen, J.S. Guiltinan, S. Beleza, J. Rocha, R.W. Pereira, and M.D. Shriver. Iris texture traits show associations with iris color and genomic ancestry. *American Journal of Human Biology*.

[27] J.E. Richman, K. Golden McAndrew, D. Decker, and S.C. Mullaney. An evaluation of pupil size standards used by police officers for detecting drug impairment. *Optometry-Journal of the American Optometric Association*, 75(3):175–182, 2004.

[28] J.M. Seddon, C.R. Sahagian, R.J. Glynn, RD Sperduto, and ES Gragoudas. Evaluation of an iris color classification system. the eye disorders case-control study group. *Investigative ophthalmology & visual science*, 31(8):1592–1598, 1990.

[29] J. Shlens. A tutorial on principal component analysis. *Measurement*, 51, 2005.

[30] J. Sim and C.C. Wright. The kappa statistic in reliability studies: use, interpretation, and sample size requirements. *Physical therapy*, 85(3):257–268, 2005.

[31] J. Sivic, B.C. Russell, A.A. Efros, A. Zisserman, and W.T. Freeman. Discovering object categories in image collections. 2005.

[32] R.A. Sturm and M. Larsson. Genetics of human iris colour and patterns. *Pigment cell & melanoma research*, 22(5):544–562, 2009.

[33] E. Tola, V. Lepetit, and P. Fua. A fast local descriptor for dense matching. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. Ieee, 2008.

[34] E. Tola, V. Lepetit, and P. Fua. Daisy: An efficient dense descriptor applied to wide-baseline stereo. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(5):815–830, 2010.

[35] S. Walsh, F. Liu, K.N. Ballantyne, M. van Oven, O. Lao, and M. Kayser. Irisplex: a sensitive dna tool for accurate prediction of blue and brown eye colour in the absence of ancestry information. *Forensic Science International: Genetics*, 2010.