

Iris Imaging for Health Diagnostics

by

Tania Weidan Yu

S.B. E.E.C.S., Massachusetts Institute of Technology, 2017

Submitted to the

Department of Electrical Engineering and Computer Science
in Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

Massachusetts Institute of Technology

June 2018

© Massachusetts Institute of Technology 2018. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 25, 2018

Certified by
Richard Fletcher
Research Scientist, Thesis Supervisor
May 25, 2018

Accepted by
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Iris Imaging for Health Diagnostics

by

Tania Weidan Yu

Submitted to the Department of Electrical Engineering and Computer Science
on May 25, 2018, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

The development of mobile technology and machine learning tools has made it easier than ever to monitor health without visiting a doctor. In this thesis, we explore the use of iris imaging as a medical diagnostic tool. We implement a system in which images captured using a mobile device can be uploaded to and analyzed by a central server. With this platform, we hope to build a large database of standard iris images with labeled medical data and facilitate studies of iris diagnostics. In our implementation, the feature extraction and classification tools built are applied to predict diabetes, through a study conducted in collaboration with researchers at Swami Vivekananda Yoga Anusandhana Samsthana (SVYASA). The results show improvement in prediction accuracy and encourage further development of the server platform for future, large-scale studies.

Thesis Supervisor: Richard Fletcher
Title: Research Scientist

Acknowledgments

I would like to thank my advisor, Dr. Richard Fletcher, for his guidance and mentorship along every step of the project. I learned a lot this past year, and I am grateful to have had the opportunity to work in this lab. I would also like to thank all my loved ones who supported me throughout the process, making the bad times good and the good times better.

Contents

1	Motivation: Non-Invasive Health Diagnostics	11
1.1	Need for Ubiquitous Health Monitoring	11
1.2	Limited Healthcare Facilities in Developing Countries	12
1.3	Work in the MIT Mobile Technology Group	12
1.4	Eye Imaging for Diagnostics	13
2	Background: Iris Physiology	15
2.1	Anatomy of the Iris	15
2.2	Development of the Iris	16
3	Related Work	19
3.1	Iris Biometrics	19
3.2	Retinal Diagnostics	20
3.3	Iridology	20
4	Proposed Solution: Iris Diagnostics Platform	23
4.1	Motivation	23
4.2	System Architecture	24
5	Feature Extraction for Iris Analysis	25
5.1	Iris Segmentation	25
5.2	Texture Features	27
5.2.1	Gabor Features	27
5.2.2	Other Features	32

5.3	Color Features	33
6	Machine Learning Algorithms for Iris Diagnostic Prediction	35
6.1	Linear Classifiers	35
6.2	Non-Linear Classifiers	36
7	Implementation of Iris Analysis for Diabetes Screening	39
7.1	Data Collection	39
7.1.1	Imaging Device	39
7.1.2	Mobile Application	40
7.1.3	File Server	41
7.2	Data Analysis	41
7.3	Results	43
7.4	Discussion	45
8	Conclusion and Future Work	47

List of Figures

2-1	Diagram of the lower part of the right eye[7].	16
2-2	Extremely close-up image of an eye that shows the fibrous structure on the anterior border layer and stroma[13].	17
3-1	Iridology chart of the left (top) and right (bottom) eye irises [23] . . .	22
4-1	System architecture for iris diagnostics platform	24
5-1	The black circle marks the detected iris, and the white circle marks the detected pupil.	26
5-2	Sampling from ROI polar coordinates to transform into rectangular region	27
5-3	Gabor filters in the spatial dimension, for $\theta = \{0, \frac{\pi}{8}, \dots, \frac{7\pi}{8}\}$ and $f = \{1, 2, 3, 4\}$	28
5-4	Output images as a result of applying an array of Gabor filters to a 5-by-5 checkerboard. The filters used have frequencies $\{0.01, 0.02, 0.03\}$ and orientations $\{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}\}$	29
5-5	Plot of pixel variance v. frequency of Gabor filter for two checkerboards: 10-by-10 and 20-by-20. The orientation is $\theta = 0$ for all filters applied. Note that the variances for the 20-by-20 checkerboard are scaled by a constant so that the peaks can be more easily compared. .	30
5-6	Heatmap of Gabor response variances. The x axis values represent frequencies in the range $[0.1, 2]$, and the y axis values represent angles in the range $[0, \frac{7\pi}{8}]$	31

7-1	The imaging device consists of a webcam with six infrared LEDs, as well as a hood attachment that covers the webcam. The device is connected to a mobile device which is running the companion mobile application.	40
7-2	The iris is labeled with 24 ROI, each corresponding to a different organ, system, or part of the human body. The black circle in the center is the pupil. The regions are reflected along the vertical axis in the left eye.	42
7-3	The image on the left shows an eye in which part of the iris is occluded by the upper and lower eyelids. The image on the right shows an eye in which the iris is completely visible.	42
7-4	Mean ROC curve for 3-fold cross validation using a linear SVM model with LGBPHS features.	44

Chapter 1

Motivation: Non-Invasive Health Diagnostics

1.1 Need for Ubiquitous Health Monitoring

Traditionally, healthcare has been reactive; when patients notice symptoms of being ill, they contact a doctor and seek treatment. In recent years, the healthcare industry in countries such as the United States and Canada has been shifting toward a preventive care model, in which healthcare systems intervene much earlier, upon detection of risk and even attempting to predict risk[21]. The social benefits are clear, as early intervention means that people are treated before the onset of symptoms even occurs. This can also reduce the economic burden of disease, both in terms of treatment costs as well as loss of productivity[14]. Preventive care is particularly important because of the aging population across the world. By 2030, the estimated population over the age of 60 is nearly 1.5 billion, as opposed to the just over 900 million in 2015[16]. If we were able to improve an individual's overall health, he or she would suffer fewer ailments with age and incur less cost on the healthcare system.

An important enabler of this change is technological development. With the widespread access to smart phones and other Internet-connected devices, health monitoring has become increasingly accessible and powerful. In 2017, a study published by Aruba reported that 60% of healthcare organizations across the globe have incor-

porated Internet of Things (IoT) technology into their work[3]. This includes devices such as wearables that analyze vitals, home monitoring systems that alert caretakers when an elderly person has fallen, and even a digital mirror that measures heart rate[26]. Such automated systems, especially passive, non-contact ones, have the potential of making healthcare a seamless portion of people's lives.

1.2 Limited Healthcare Facilities in Developing Countries

In many developing countries, access to healthcare facilities and medical professionals remains limited. The lack of access to diagnostic facilities accounts for a significant portion of untreated major illnesses in rural India in particular[5]. More than 70% of the population lives in rural areas, but urban areas have more than two-thirds of the government hospital beds and a much greater private sector presence[4]. This means that many people may not be willing to seek out medical attention until they are quite ill. Various factors contribute to this issue, including lack of roads and communication infrastructure, lack of fully trained professionals, and inadequate public health financing[4].

1.3 Work in the MIT Mobile Technology Group

In MIT D-Lab, Dr. Richard Fletcher is leading the Mobile Technology Group in developing non-invasive and low-cost mobile health tools to solve some of these above-mentioned issues. Current and past projects include using a peak flow meter to detect pulmonary diseases, mobile games for monitoring mental health, and thermal imaging for diagnosis of infectious diseases. An important focus of the group's work is mobility. In developing areas in particular, a portable tool has a much wider reach, since it can be deployed to local practitioners and even individuals for usage.

1.4 Eye Imaging for Diagnostics

We are interested in the exploration of iris imaging as another diagnostic tool. In certain branches of non-Western medicine, the iris is thought to reflect conditions of the human body. Specifically, iridology is the systemic diagnosis of human bodily conditions through observed changes in the iris[17]. Most commonly, iridologists assert that disease or other medical conditions can cause abnormalities in pigmentation[11]. The method has largely remained untested – only a handful of controlled, scientific studies have been conducted to evaluate its effectiveness, and those that do exist have been inconclusive[11]. If proven effective, however, it could be a useful screening tool to aid physicians and patients alike in deciding whether or not to continue with further testing and treatment, since it requires little more than a small camera.

In Chapter 2 of this thesis, we provide an overview of iris physiology. Chapter 3 covers existing literature of various forms of automated eye image classification, as well as iridology. We then present our proposed iris diagnostics platform in Chapter 4 and the feature extraction tools in Chapter 5. Chapter 6 explains the machine learning algorithms explored in this study. Additionally, we describe our implementation for the specific application of diabetes screening and discuss our results in Chapter 7. Finally, in Chapter 8 we explain the conclusions drawn from our study along with proposals for future work on this project.

Chapter 2

Background: Iris Physiology

2.1 Anatomy of the Iris

The eye is composed of three layers: an outer layer of connective tissue, a middle vascular layer, and an inner neural layer. The main structures can be seen in Figure 2.1, which depicts the lower part of the right eye. The outer layer includes the cornea, which is transparent, and the sclera, commonly known as the "white" of the eye. The inner layer, or the retina, contains photoreceptors, which transform photon energy into neural signals. The middle layer is known as the uvea and consists of the iris, ciliary body, and choroid.[29]

The iris, a colored ring of muscles and connective tissue, is the outermost of the three uvea structures and controls dilation of the pupil, which varies the amount of light that passes through. The radius of the iris is 12 mm on average, and the radius of the pupil can range from 0.5 mm to 4.5 mm depending on dilation. A circular ridge, the collarette, lies approximately 1.5 mm from the pupillary margin and divides the iris into the pupillary zone and the ciliary zone. The ciliary zone, which is the outer ring of the iris, extends to the iris root.[29] Figure 2.1 clearly depicts the ridge formed by the collarette.

The iris can be further divided into four layers: the anterior border layer, stroma and sphincter muscle, anterior epithelium and dilator muscle, and posterior epithelium. The anterior border layer is an interweaving network of pigmented melanocytes,




Figure 2-1: Diagram of the lower part of the right eye[7].

fibroblasts, and collagen fibrils. The stroma also contains many pigmented and non-pigmented cells, along with collagen fibrils and ground substance. Arteries branch off the major circle of the iris, located in the ciliary body, and extend toward the pupil margin. The structures and density of the components in the anterior border layer and stroma, as well as pigment density in the melanocyte, contribute to the color of the iris. The sphincter muscle is a circular muscle which causes constriction of the pupil in miosis, while the dilator muscle causes dilation in mydriasis. The anterior and posterior epithelium include pigmented epithelium, and in particular part of the posterior epithelium forms the pupillary ruff, a roughly textured ring encircling the pupil.[29]

2.2 Development of the Iris

During embryonic development, the neural tube ectoderm and mesenchyme eventually give rise to the iris. After week 3, the optic cup forms as the precursor to the eye. The anterior portion of the optic cup forms the iris and ciliary body epithelia. The anterior portion of the optic cup's inner layer, most of which forms the neural



Figure 2-2: Extremely close-up image of an eye that shows the fibrous structure on the anterior border layer and stroma[13].

retina, gives rise to the non-pigmented cells of the iris. Neural crest cells form the stroma, and the sphincter and dilator muscles develop from neuroectoderm. Most of this development occurs before week 10 of the embryo's life.[25]

Chapter 3

Related Work

Eye image classification is used extensively in medicine and in biometrics, and there has been a recent focus in iris biometrics and retinal diagnostics in particular. Although there is less literature regarding iridology, there is also interest in the evaluation of iris diagnostics.

3.1 Iris Biometrics

Most applications of iris-related imaging have been designed for biometrics and identification of individuals. One of the earliest conceptual designs for automated iris biometrics was patented in 1987 by Flom and Safir, but was unimplemented[6]. In the 1990s, Daugman began laying the groundwork for much of iris biometrics fundamentals, and in fact much of commercial technology is still based on his work[6].

Because of the interest in applying iris biometrics to security, one focus has been in developing systems that can detect and analyze the iris from a distance and while the individual is in motion. Matey et al. presented the IOM system in 2006, a commercial system that can perform iris recognition while the individual walks through a confined space at normal speed[22]. In 2014, Tan and Kumar used Zernike moments to match iris features, in an effort to develop a system that can accurately recognize irises in both near-infrared and visible illumination, and in unconstrained environments[32]. The particular challenges in these unconstrained environments is achieving sufficient

image quality and handling motion blur. Daugman suggests that a minimum of 70 pixels is necessary in the iris radius, which means distant imaging systems have high resolution requirements[9]. On the other end of the spectrum, Hariprasath and Mohan have focused on personal identification systems instead, optimizing for close-up iris recognition[15].

3.2 Retinal Diagnostics

A part of the eye that is commonly used in imaging-based diagnostics is the retina. Located in the back of the eye, the retina can be imaged via a fundus camera, ophthalmoscope, or optical coherence tomography, to name a few[20]. Fundus imaging is most commonly used as a screening tool for diabetic retinopathy, a disease stemming from diabetes that damages the eye, even leading to blindness[20]. A variety of other chronic illnesses can also lead to complications of the eye, especially cardiovascular disease and neurological disease pertaining the central nervous system[20]. In 2016, Varun et al. applied deep learning algorithms to detect diabetic retinopathy with high sensitivity and specificity[34]. More recently, researchers at Google conducted a study in which deep learning of retinal images successfully predicts a variety of cardiovascular risk factors, such as age, gender, and blood pressure[27]. There is particular interest in pairing these AI algorithms with fundus cameras, since it is simple to attach an ophthalmoscopic lens to a smartphone camera, in effect creating a fundus camera optical system, and capture good quality images[18].

3.3 Iridology

In iridology, the key idea is that specific regions of the iris map to various organs and tissues in the human body [17]. We can see in Figure 3.3 a chart developed by Bernard Jensen. The chart divides the left and right eye irises into several rings and further splits them into sections based on degrees around the circle.

In the past, studies of iridology have shown inconclusive results, and the field has

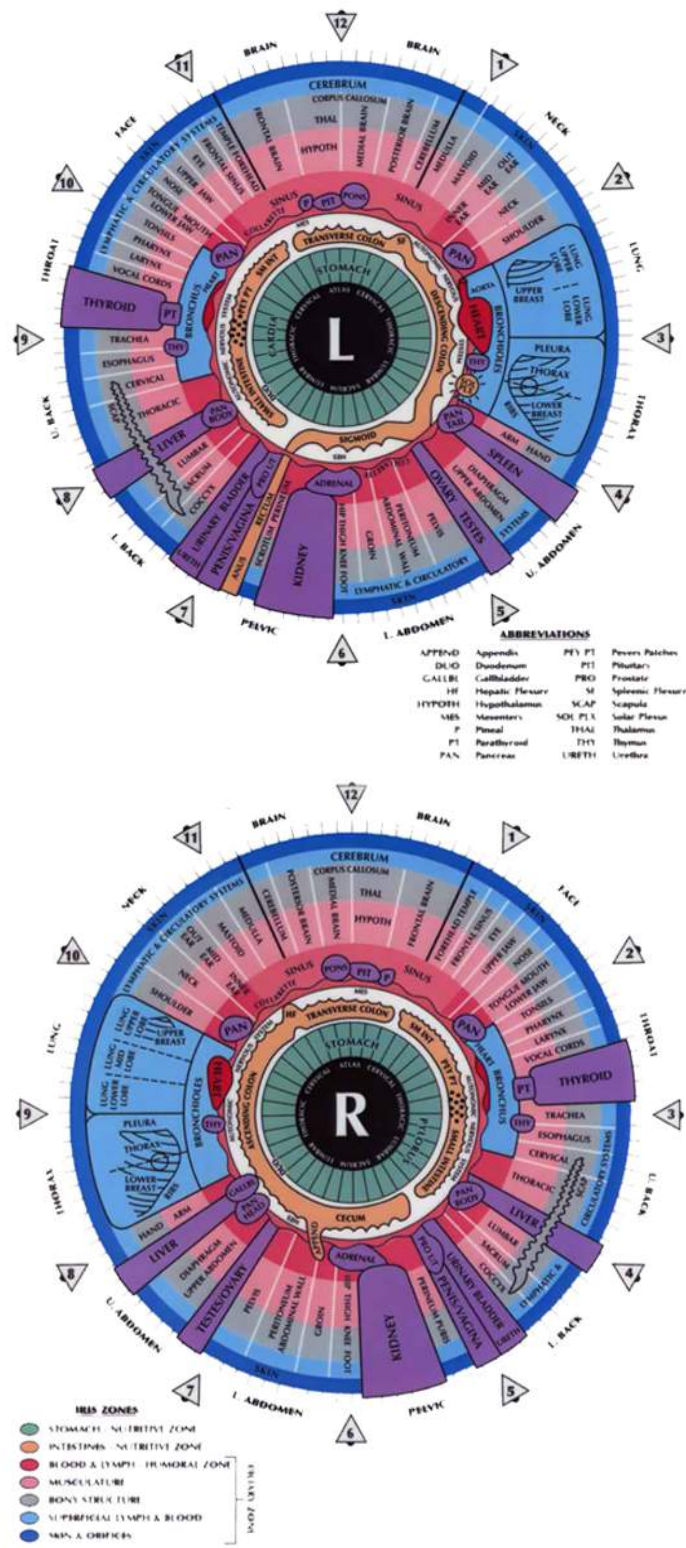


Figure 3-1: Iridology chart of the left (top) and right (bottom) eye irises [23]

not been conventionally accepted in western medicine. A 2000 review of iridology studies found only four controlled and masked experiments on the subject, and those either return results similar to random chance or have too-small sample sizes[11]. In 2007, however, Discant et al. automate the process to more successfully classify iris images from healthy individuals and those with cardiovascular problems, using b-spline wavelets and moments from the images [10]. A recent study performed by researchers at Swami Vivekananda Yoga Anusandhana Samsthana (SVYASA) and Ramaiah University also shows potential in this method of diagnosis for diabetes [33].

Chapter 4

Proposed Solution: Iris Diagnostics Platform

4.1 Motivation

As mentioned in Section 3.3, there have been few attempts at conducting scientific studies on the effectiveness of iris diagnostics. With the recent advent in machine learning techniques and IOT (Internet of Things) technology, we now have the tools to create a large-scale platform for storage and analysis of iris images, in conjunction with other factors of an individual's health. Using these, we would be able to obtain more meaningful and conclusive results than previous studies have managed.

We propose the development of an online platform and mobile application that enables easy upload and analysis of iris images. Any individual may download the application, capture an image of their iris, and receive results within minutes. We also intend to provide access to the data collected to any individuals interested in performing their own analyses on the raw information. Of course, due to the sensitive nature of the data (iris biometrics could allow identification of the study participants), this access will be limited to those explicitly given permission. By building a large database of iris images labeled with multiple health factors and providing tools for basic feature extraction, we aim to not only perform our own analyses but also encourage others to participate as well.

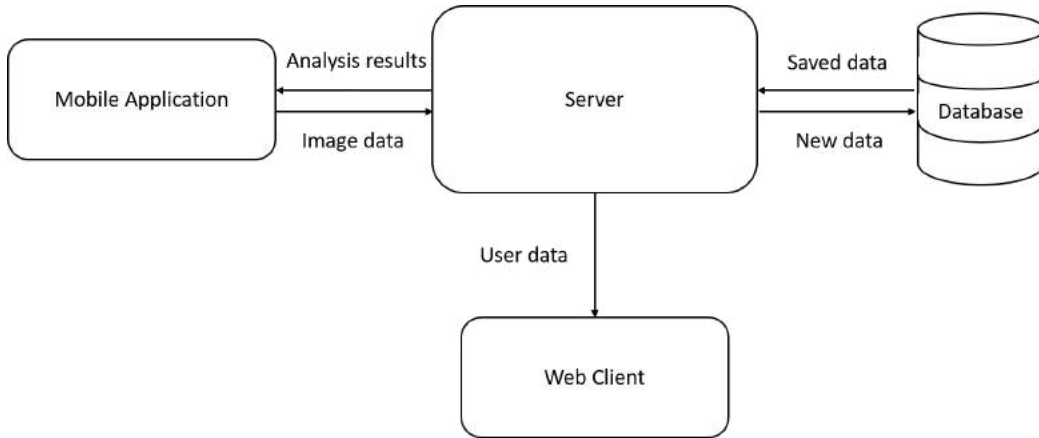


Figure 4-1: System architecture for iris diagnostics platform

4.2 System Architecture

The system provides two services: analysis of personal iris data and access to overall study results and data. The tasks are divided among three main components: a mobile application, web server, and frontend web UI, as depicted in Figure 4.2. The user captures an image of the eye using either the phone camera or an attached imaging device and uploads it via the application, along with other basic health information such as age, weight, or blood pressure. Personal analysis results can also be viewed in the application upon log in. The server stores and processes iris images, extracting the appropriate features and calculating scores for each image. These features and scores are stored within the database on the server as well. The web interface allows access to the raw feature data with proper permissions, along with general statistics about and visualizations of the data.

Chapter 5

Feature Extraction for Iris Analysis

To perform feature extraction, we must first preprocess and segment the target region of the iris. We then examine a range of features, which include texture features to detect structural changes as well as color features for pigmentation changes.

5.1 Iris Segmentation

We prepare for iris segmentation by enhancing the grayscale images in two ways: increasing contrast and removing glare. After experimentation with different parameters, we find that limiting the grayscale pixel range to $[0.15, 0.7]$ sufficiently enhances the image to correctly detect the iris. To remove glare, we use morphological reconstruction by dilation, a geodesic operator that fills "holes" in an image. This process involves the original image I and a structuring element S , which defines a neighborhood surrounding any single pixel. We define gray-valued dilation as

$$\delta(I)(x) = I(\max_{s \in S} x + s) \tag{5.1}$$

where x and s are vector quantities[1]. Conceptually, each pixel takes on the maximum value in its neighborhood. When this dilation process is constrained beneath a mask of the original image, the end result is a copy of the original image with its regions of lower value filled.

After preprocessing is complete, we find the position and size of the iris. Many such methods exist, and they are generally variations of circle detection in an image. One of the most common ones performs a search using the Daugman integrodifferential operator [8]

$$\max_{r, x_0, y_0} G_\sigma(r) * \frac{\partial}{\partial r} \oint_{r, x_0, y_0} \frac{I(x, y)}{2\pi r} ds \quad (5.2)$$

where $I(x, y)$ is the image, r is the radius, and (x_0, y_0) are the center coordinates. $G_\sigma(r)$ is a smoothing function with scale σ . By finding the circle with the greatest contour integral derivative, we can detect the edge of the pupil or iris, depending on the range of coordinates and radii over which the algorithm searches. Figure 5.1 shows the circles found to delineate the iris and pupil. Because of the fine search required to detect the iris and pupil, this portion of the process is quite slow. We increase the speed by searching in a scaled down copy of the image for a coarse estimate, then searching within only a small range of the estimated center and radius in the original image.




Figure 5-1: The black circle marks the detected iris, and the white circle marks the detected pupil.

After finding the iris, it is simple to locate a particular region of interest (ROI), which is specified by a range over angle and a range over distance from the center.

In order to convert this curved region into a rectangle, we sample from its polar coordinates for easier analysis, as seen in Figure 5.1.




Figure 5-2: Sampling from ROI polar coordinates to transform into rectangular region

5.2 Texture Features

An important component of an image is its texture: the presence of patterns, edges, or other features that are unrelated to color. A class of such methods is wavelets. One commonly used wavelet in signal processing is the Gabor wavelet, and its 2-D sister can be used for image texture extraction. Other methods of capturing texture include applying Fourier transform and calculating the gray-level co-occurrence matrix (GLCM).

5.2.1 Gabor Features

Gabor Filter

The 2-dimensional Gabor filter is a modulation of a Gaussian function with a complex sinusoidal, allowing it to capture both orientation and frequency characteristics [31].




Figure 5-3: Gabor filters in the spatial dimension, for $\theta = \{0, \frac{\pi}{8}, \dots, \frac{7\pi}{8}\}$ and $f = \{1, 2, 3, 4\}$.

It is represented by the following equation:

$$G(x, y) = e^{-\frac{1}{2}\left(\frac{x'^2}{\sigma_x^2} + \frac{y'^2}{\sigma_y^2}\right)} e^{i(2\pi x' f + \psi)} \quad (5.3)$$

$$x' = x \cos \theta + y \sin \theta \quad (5.4)$$

$$y' = -x \sin \theta + y \cos \theta \quad (5.5)$$

where (x, y) is a pixel position in the image, σ_x and σ_y are the standard deviations of the Gaussian in the x and y directions, f is the frequency of the sinusoid, ψ represents the phase offset, and θ is the orientation. Figure 5.2.1 depicts filters of varying orientations and frequencies which are used in this study. We use these filters to capture a comprehensive set of features for an image I by taking the output of the the convolution $G * I$, for each filter G in a particular orientation θ and frequency f .

By specifying a number of Gabor filters in various directions and frequencies,




Figure 5-4: Output images as a result of applying an array of Gabor filters to a 5-by-5 checkerboard. The filters used have frequencies $\{0.01, 0.02, 0.03\}$ and orientations $\{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}\}$.

we can capture a comprehensive set of features. We can see in Figure 5.2.1 the effects of applying an array of Gabor filters to a checkerboard image. For filters of orientation $\theta = k\pi$, for $k \in \mathbb{Z}$, vertical lines are highlighted in the output; for filters of orientation $\theta = \frac{k\pi}{2}$, horizontal lines are highlighted. More concretely, we expect that when a Gabor filter captures relevant feature information, the variance of output pixels will be higher. Figure 5.2.1 plots the pixel variance across frequencies for two checkerboards. The variance peaks at $f = 0.05$ for the 10-by-10 checkerboard, and $f = 0.1$ for the 20-by-20 checkerboard. This is consistent with expectation, since the 20-by-20 checkerboard has twice the base frequency as does the 10-by-10.

When we apply this feature extraction method to iris images, we perform similar frequency and orientation analyses to select the most useful features. Figure 5.2.1 shows a heatmap of where the variance peaks for a sample section of the iris. The

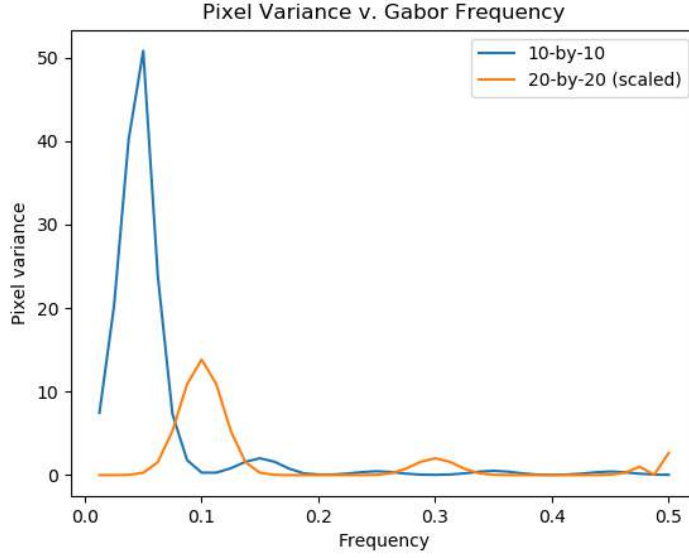


Figure 5-5: Plot of pixel variance v. frequency of Gabor filter for two checkerboards: 10-by-10 and 20-by-20. The orientation is $\theta = 0$ for all filters applied. Note that the variances for the 20-by-20 checkerboard are scaled by a constant so that the peaks can be more easily compared.

peaks correspond to the (f, θ) pairs: $(0.98, 0)$, $(0.98, \frac{\pi}{2})$, $(1.44, \frac{\pi}{4})$, and $(1.44, \frac{3\pi}{4})$.

Local Gabor Binary Pattern Histogram Sequence

One common method of extracting features from Gabor filter responses is applying the local binary pattern (LBP) operator[35]. The LBP operator thresholds each of the 8 neighbors of a pixel to a 0 or 1, such that

$$S(f_p - f_c) = \begin{cases} 1 & \text{if } f_p > f_c \\ 0 & \text{if } f_p < f_c \end{cases} \quad (5.6)$$

The LBP pattern at each pixel is then calculated as

$$LBP = \sum_{p=0}^7 S(f_p - f_c) 2^p \quad (5.7)$$




Figure 5-6: Heatmap of Gabor response variances. The x axis values represent frequencies in the range $[0.1, 2]$, and the y axis values represent angles in the range $[0, \frac{7\pi}{8}]$.

Since we apply LBP to the Gabor magnitude response, this returns the local Gabor binary pattern (LGBP).

To reduce the feature space and to find regional properties of these patterns, we additionally find histogram sequences of the LGBP (LGBPHS). The image is divided into non-overlapping patches, and gray-level histograms are extracted from each patch.

5.2.2 Other Features

Fast Fourier Transform

In addition to Gabor wavelets, we also apply a 2-dimensional fast Fourier transform (FFT) to the image, which breaks the image down into sinusoidal components. This allows us to capture the overall frequency information but without specific orientation.

Gray-Level Co-Occurrence Matrix

Proposed by R. M. Haralick, the gray-level co-occurrence matrix (GLCM) is a measure of spatial relationship between pixels [28]. A GLCM C of size $k \times k$ is defined such that C_{ij} is the frequency at which a pixel of gray level value g_i occurs some specified vector t away from a pixel of value g_j , where there are k distinct gray levels. Using this, we can calculate a number of features:

1. Energy

$$\sum_{i=1}^k \sum_{j=1}^k C_{ij}^2 \quad (5.8)$$

2. Contrast

$$\sum_{i=1}^k \sum_{j=1}^k C_{ij} (i - j)^2 \quad (5.9)$$

3. Entropy

$$-\sum_{i=1}^k \sum_{j=1}^k C_{ij} \log C_{ij} \quad (5.10)$$

4. Correlation

$$\sum_{i=1}^k \sum_{j=1}^k \frac{C_{ij}(i - \mu_i)(j - \mu_j)}{\sigma_i \sigma_j} \quad (5.11)$$

5. Homogeneity

$$\sum_{i=1}^k \sum_{j=1}^k \frac{C_{ij}}{1 + |i - j|} \quad (5.12)$$

The GLCM we calculate is based on a range of angles $\{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}\}$ and a vector distance of 1.

5.3 Color Features

One of the leading theories regarding iris diagnostics is that the changes are reflected in pigmentation. Therefore, for the RGB images collected, we want to examine features within the colorspace as well. One commonly used feature is color histograms. We divide each of the three color channels into bins, and each bin stores the number of pixels in the image with corresponding RGB values.

Chapter 6

Machine Learning Algorithms for Iris Diagnostic Prediction

We explore a variety of classification methods that might be used for diagnostic prediction. In its most basic form, diagnosis is a binary classification problem.

6.1 Linear Classifiers

One simple classifier is logistic regression, which estimates the probability of a label y by calculating the distance from the feature vector to a hyperplane with normal θ . More specifically, we have

$$P(y = 1|x) = \frac{1}{1 + e^{-\theta^T x}} = \sigma(\theta^T x) \quad (6.1)$$

$$P(y = 0|x) = 1 - P(y = 1|x) \quad (6.2)$$

with a loss function

$$Loss = \frac{1}{n} \sum_i (y^{(i)} \log P(y^{(i)} = 1|x^{(i)}) + (1 - y^{(i)}) \log (1 - P(y^{(i)} = 1|x^{(i)}))) \quad (6.3)$$

where n is the total number of training points, $y^{(i)}$ is the label for $x^{(i)}$, the i th data point in the training data set, and θ is the parameters we train[19].

Another linear classifier is the support vector machine (SVM). We again use a linear separator θ , but instead of applying the sigmoid function to obtain a probability, we find the hyperplane that maximizes the margin. In other words, we find the hyperplane such that the distance between it and the nearest datapoints is maximized. This requires the use of hinge loss:

$$Loss = \frac{1}{n} \sum_i \max\{0, 1 - y^{(i)} \theta^T x^{(i)}\} \quad (6.4)$$

In both SVM and logistic regression, we can additionally include a regularization term $\lambda \|\theta\|^2$ in the loss function, where λ controls the strength of regularization.

6.2 Non-Linear Classifiers

SVM is often enhanced by applying different kernels to the data points, which allows us to find non-linear separators even though it is inherently a linear model. One such kernel is the radial basis function (RBF). This applies a Gaussian to each point in the training data set to form a basis and calculates features based on the proximity of each data point to this basis[19].

In this study, we also examine a class of non-linear models, decision trees. At each node of a tree, the data is split based on a condition, usually on a single feature. The leaves of the tree contain label predictions. Many algorithms exist for the construction of decision trees, most using a top-down approach and picking split conditions at each node that will maximize some criterion. Two of the most commonly used criteria for splitting are entropy and Gini index[30]. Change in entropy, or information gain, represents the amount of information learned about the labels from performing a split

$$H(S) - \sum_{s \in \{S_l, S_r\}} \frac{|s|}{|S|} \cdot H(s) \quad (6.5)$$

where

$$H(S) = \sum_{y \in \{0,1\}} -\frac{|S_y|}{|S|} \log_2 \frac{|S_y|}{|S|} \quad (6.6)$$

such that H is entropy, S is the total set of data points, S_l and S_r are the subsets contained within the left and right child nodes, and S_y is the subset of data points with label y . The criterion based on the Gini index has a similar form

$$G(S) = \sum_{s \in \{S_l, S_r\}} \frac{|s|}{|S|} \cdot G(s) \quad (6.7)$$

where

$$G(S) = 1 - \sum_{y \in \{0,1\}} \left(\frac{|S_y|}{|S|} \right)^2 \quad (6.8)$$

where G is the equation for the Gini index. For both criteria, a higher information gain or Gini gain means that the impurity of the data has been reduced and is therefore more desirable.

Chapter 7

Implementation of Iris Analysis for Diabetes Screening

To evaluate the viability of our proposed platform, we implement a basic version of the system and apply it to screening of diabetes. We conduct this study in collaboration with researchers at the diabetes clinic of SVYASA in Bengaluru, India. Diabetes is a widespread chronic condition that can lead to many other complications if not diagnosed and treated early[2]. In 2014, 422 million people across the world were diagnosed with diabetes [24]; in India, there were close to 70 million [12]. Current methods of diagnosis require performing blood tests on the patient, which is invasive and inconvenient for those who do not have ready access to healthcare professionals. Therefore, successful application of iris diagnostics to diabetes could be very impactful.

7.1 Data Collection

7.1.1 Imaging Device

To capture images of the iris, we use a webcam connected to a mobile device, as pictured in Figure 7.1.1, or the camera on the device itself. The camera is optionally surrounded by six infrared LED lights in order to eliminate the effects of different



Figure 7-1: The imaging device consists of a webcam with six infrared LEDs, as well as a hood attachment that covers the webcam. The device is connected to a mobile device which is running the companion mobile application.

iris pigmentations. This also includes a 3D-printed hood attachment, designed by another member of the group, that covers the eye, which allows for consistent camera distance. It also limits the effects of surrounding light, such as glare, especially when infrared illumination is used.

7.1.2 Mobile Application

The mobile phone or tablet connected to the imaging device contains an application that controls image capture and data transmission. Once the image of the iris is captured, the user can send it to the server for analysis, along with responses to a general health questionnaire, which may include information such as eye color, ethnicity, diet, or current medication. Once the server completes its feature analysis, the results are displayed in the application. Presently, these are scores for each region in the iris.

7.1.3 File Server

The server framework is developed by another member of the group, using Django and a PostgreSQL database. The machine learning work is implemented in Python.

When an image is sent to the server, it is processed for feature extraction and analysis. We detect the iris as described in Section 5.1 and segment each desired region as a rectangular matrix. We then extract a feature vector from this region. These features are compared to the existing features in our database, and scores are computed for each region, calculated as a distance metric from the average individual's iris. This can be generalized, however, to other metrics such as prediction score for a particular disease. The resulting features are also saved to the database, so that our dataset is continuously growing.

For the particular application of diabetes screening, we also implement the classification of images in a labeled dataset. This uses the algorithms described in Chapter 6 to predict the extracted feature vectors as diabetic or non-diabetic. Because our dataset is limited, we elect to use relatively simple classification tools in order to avoid unduly complex models and overfitting to the data.

7.2 Data Analysis

Our dataset consists of 11 iris images, collected from 8 study participants at the SVYASA diabetes clinic. Using the iridology chart in Figure 3.3, we divide each iris into 24 regions, as depicted in Figure 7.2. If we compare these regions to a sample iris image in Figure 7.2, we can see that some regions are partially or entirely occluded by the upper and lower eyelids. During segmentation, we detect whether the target region is completely visible by performing a simple frequency analysis using FFT. Because the skin is mostly uniform, whereas the iris has very detailed texture, regions that are partly or entirely covered have lower high frequency coefficients. If a region is not wholly visible, we ignore it in our analysis. For the application to diabetes screening, we select region 21, which corresponds to the pancreas, the organ that produces insulin and regulates blood glucose levels.




Figure 7-2: The iris is labeled with 24 ROI, each corresponding to a different organ, system, or part of the human body. The black circle in the center is the pupil. The regions are reflected along the vertical axis in the left eye.



Figure 7-3: The image on the left shows an eye in which part of the iris is occluded by the upper and lower eyelids. The image on the right shows an eye in which the iris is completely visible.

Feature	LGBPHS	Gabor response	Color histogram	GLCM	FFT
Average time (s)	18.7	14.9	4.5	0.84	0.13
Dimensions	7200	2457600	512	180	307200

Table 7.1: Average time in seconds and feature dimension for extracting each feature type from a region of the iris.

The average iris radius is 190 pixels across, and the average pupil radius 52 pixels. Iris detection is one of the most time-consuming processes in the pipeline, requiring about 18 seconds per image. Table 7.2 lists the average time required for each type of feature extraction performed on a single region. Raw Gabor response is by far the largest feature, since it is essentially multiple filtered versions of the original ROI. We also see that, due to the complexity of the Gabor wavelet, applying the Gabor filter is computationally more consuming than all other features. FFT is the fastest feature to calculate, although its size is also quite large.

7.3 Results

We evaluate the performance of our classification models on two metrics: percent accuracy and area under the receiver operating characteristic (ROC) curve. The ROC curve plots the true positive rate against the false positive rate at various prediction thresholds, where a perfect classifier would have an area under the curve (AUC) of 1. Table 7.3 presents the average accuracy and AUC results from 3-fold cross-validation on a variety of region 21 feature and machine learning model combinations.

To evaluate our hypothesis that different regions of the iris reflect the condition of different parts of the body, we perform classification on other regions as well. Because a portion of the regions are partially occluded in some images, and some are affected by glare, we include in Table 7.3 only the results from regions that are visible and without glare in all 11 images.

Algorithm	LGBPHS		Gabor response		Color histogram		GLCM		FFT	
	Acc	AUC	Acc	AUC	Acc	AUC	Acc	AUC	Acc	AUC
LR	0.76	0.89	0.53	0.67	0.64	0.61	0.58	0.39	0.76	0.94
Linear SVM	0.76	0.89	0.64	0.75	0.64	0.45	0.53	0.34	0.76	0.89
RBF SVM	0.64	0.83	0.64	0.5	0.64	0.5	0.64	0.5	0.64	0.61
Decision tree	0.7	0.66	0.57	0.57	0.74	0.66	0.93	0.91	0.55	0.56

Table 7.2: Mean accuracy and AUC scores from running 3-fold cross-validation on each feature and algorithm pair, for region 21. The algorithms used are logistic regression (LR), linear SVM, RBF-kernel SVM with $\gamma = 0.0001$, and decision tree using Gini index. The features applied are LGBPHS, raw Gabor responses, color histograms, GLCM properties, and FFT coefficients.

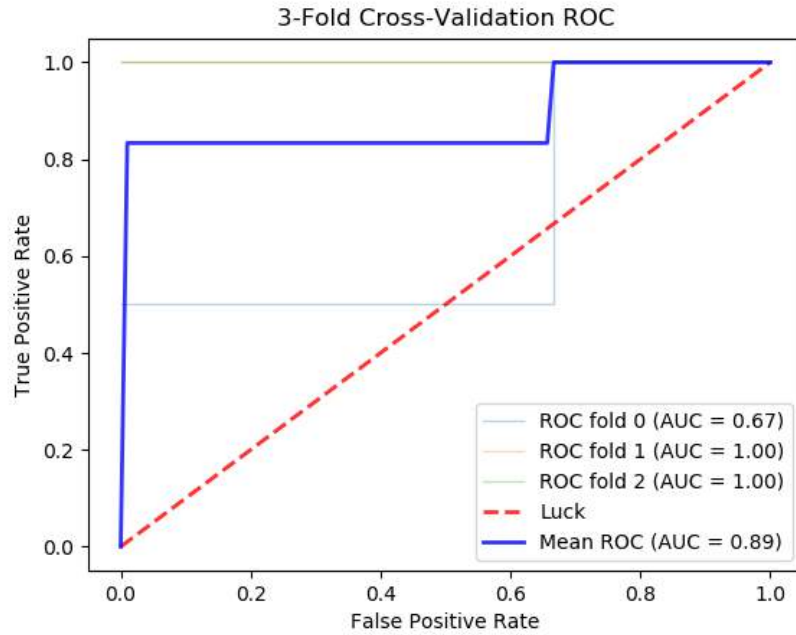


Figure 7-4: Mean ROC curve for 3-fold cross validation using a linear SVM model with LGBPHS features.

Region	6	15	18	21	23	24
Accuracy	0.64	0.58	0.76	0.76	0.6	0.47
AUC	0.42	0.72	0.99	0.89	0.28	0.01

Table 7.3: Mean accuracy and AUC scores from running 3-fold cross-validation on a subset of iris regions. A linear SVM is used with LGBPHS features.

7.4 Discussion

We can see from results in Table 7.3 that GLCM properties return the best scores when applied to the decision tree, with accuracy of 0.93 and AUC of 0.91. However, they perform very poorly with all other models. LGBPHS and FFT have the next best scores, for the logistic regression and linear SVM models. All four cases have a 0.76 accuracy, with AUC of at least 0.89. Because these both have large feature spaces, and especially because the size of our dataset is so limited, it is reasonable that the simpler, linear classifiers would perform better than the complex, non-linear ones, which would likely overfit. Raw Gabor response and color histogram both have relatively low scores overall. Although GLCM achieves the best results, LGBPHS is most consistent across different classification algorithms. Therefore, we conjecture that LGBPHS may be more generalizable.

On the other hand, if computation time is of greater importance, we may prefer FFT or GLCM instead. GLCM is over 20 times faster to compute as LGBPHS, and FFT over 100 times faster, as seen in Table 7.2. Both feature types can achieve similar or better prediction performance, with much less time. If memory is limited, we would prefer GLCM properties over the rest because of its small feature size. LGBPHS requires 40 times the memory, and FFT over 1000 times. It may be possible, however, to prune the FFT feature set and reduce the number of dimensions by selecting only the relevant frequency ranges.

The results for color features is inconclusive. Color histograms appear to be mostly ineffective for diabetes prediction, but there are also a variety of different color features that could be calculated.

When we compare across different regions of the iris in Table 7.3, we find that region 18 performs slightly better than region 21 does, while the other regions perform rather poorly. By matching region 18 to the iridology chart, we notice that region 18 is another region that corresponds to the pancreas. The other regions tested correspond to the liver, lungs, small intestine, and stomach. We cannot make a definite claim as to the relevance of different iris regions, since they are of all different sizes, and

our analysis may be impacted by that. Furthermore, we are only able to evaluate this claim on a small subset of the iris regions. However, these results suggest that there may be some significance to distinct regions of the iris and encourage further exploration.

Chapter 8

Conclusion and Future Work

In this thesis, we establish a platform for conducting large-scale studies on iris diagnostics and explore the effectiveness of using iris imaging as a diagnostic tool, with the specific application of diabetes prediction. The results from classification are encouraging, as we perform better than the baseline of random guessing with several combinations of features and machine learning algorithms. Furthermore, classification on relevant regions based on the iridology chart has better performance than on regions corresponding to other organs or systems in the body. However, the dataset is very limited, so we are unable to draw strong conclusions. With our proposed system of allowing individuals to upload iris images to a centralized database, we lay the groundwork for a large-scale study and the creation of a diverse and well-annotated dataset.

Future work in iris diagnostics should explore additional features and algorithms. In particular, different color features should be tested to evaluate the relevance of texture versus pigmentation in iridology. If color is found to be irrelevant, it may be preferable to use infrared cameras to capture iris images. With more data, deep learning techniques should also be applied. Due to their scale- and rotation-invariant qualities, convolutional neural networks could be used to better learn features (both color and texture) from the images. Furthermore, it may be interesting to analyze images in combination with other health data, such as age or other conditions, so that changes in the iris due to these other factors can be effectively ignored in classification.

Bibliography

- [1] Mathematical morphology.
- [2] Ramachandran Ambady and Snehalatha Chamukuttan. Early diagnosis and prevention of diabetes in developing countries. *Reviews in Endocrine and Metabolic Disorders*, 9(3):193, Jul 2008.
- [3] Aruba. Iot heading for mass adoption by 2019 driven by better-than-expected business results. <http://news.arubanetworks.com/press-release/arubanetworks/iot-heading-mass-adoption-2019-driven-better-expected-business-results>, February 2017.
- [4] Yarlini Balarajan, S. Selvaraj, and S. V. Subramanian. Health care and equity in india. *Lancet*, 377:505–515, 2011.
- [5] Debasis Barik and Amit Thorat. Issues of unequal access to public health in india. *Frontiers in Public Health*, 3:245, 2015.
- [6] Kevin W. Bowyer, Karen Hollingsworth, and Patrick J. Flynn. Image understanding for iris biometrics: A survey. *Computer Vision and Image Understanding*, 110(2):281–307, may 2008.
- [7] Wikimedia Commons. File:schematic diagram of the human eye en.svg — wikipedia commons, the free media repository, 2018. [Online; accessed 25-May-2018].
- [8] J. Daugman. How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):21–30, Jan 2004.
- [9] John Daugman. How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*, 14:21–30, 2002.
- [10] Anca Discant, Simina Emerich, Eugen Lupu, and Sorina Demea. Healthy and cardiac people iris classification by use of b-spline wavelet decomposition. *International Conference on Advancements of Medicine and Health Care through Technology*, September 2017.
- [11] Ernst EE. Iridology: Not useful and potentially harmful. *Archives of Ophthalmology*, 118(1):120–121, 2000.

- [12] International Diabetes Federation. Idf sea members: India. <https://www.idf.org/our-network/regions-members/south-east-asia/members/94-india.html>, 2017.
- [13] Damien Gayle. The eyes have it: The iris pictured in remarkable detail by incredible close-up shots. <http://www.dailymail.co.uk/sciencetech/article-2246888/The-eyes-The-iris-pictured-remarkable-incredible-close-shots.html>, December 2012.
- [14] Surgeon General. Economic benefits of preventing disease. <https://www.surgeongeneral.gov/priorities/prevention/strategy/appendix1.pdf>, 2011.
- [15] S. Hariprasath and V. Mohan. Biometric personal identification based on iris recognition using complex wavelet transforms. In *2008 International Conference on Computing, Communication and Networking*, pages 1–5, Dec 2008.
- [16] William A. Haseltine. Aging populations will challenge healthcare systems all over the world. <https://www.forbes.com/sites/williamhaseltine/2018/04/02/aging-populations-will-challenge-healthcare-systems-all-over-the-world/#60b114e82cc3>, April 2018.
- [17] Bernard Jensen. *Iridology Simplified*. 1980.
- [18] Hossein Nazari Khanamiri, Austin Nakatsuka, and Jaafar El-Annan. Smartphone fundus photography. *Journal of Visualized Experiments*, (125), jul 2017.
- [19] Imran Kurt, Mevlut Ture, and A. Turhan Kurum. Comparing performances of logistic regression, classification and regression tree, and neural networks for predicting coronary artery disease. *Expert Systems with Applications*, 34(1):366 – 374, 2008.
- [20] T J MacGillivray, E Trucco, J R Cameron, B Dhillon, J G Houston, and E J R van Beek. Retinal imaging as a source of biomarkers for diagnosis, characterization and prognosis of chronic illness or long-term conditions. *The British Journal of Radiology*, 87(1040):20130832, aug 2014.
- [21] Emily MacIntosh, Nirusan Rajakulendran, Zayna Khayat, and Alexis Wise. Transforming health: Shifting from reactive to proactive and predictive care. <https://www.marsdd.com/newsand-insights/transforming-health-shifting-from-reactive-to-proactive-andpredictive-care>, March 2016.
- [22] J. R. Matey, O. Naroditsky, K. Hanna, R. Kolczynski, D. J. Lofacono, S. Mangru, M. Tinker, T. M. Zappia, and W. Y. Zhao. Iris on the move: Acquisition of images for iris recognition in less constrained environments. *Proceedings of the IEEE*, 94(11):1936–1947, Nov 2006.

- [23] Healthy You Naturally. Iridology chart. <http://www.healthyyounaturally.com/>, 2017.
- [24] World Health Organization. Diabetes fact sheet. <http://www.who.int/mediacentre/factsheets/fs312/en/>, 2017.
- [25] Victoria Ort and David Howard. Development of the eye.
- [26] Ming-Zher Poh, Daniel McDuff, and Rosalind Picard. A medical mirror for non-contact health monitoring. In *ACM SIGGRAPH 2011 Emerging Technologies*, SIGGRAPH '11, pages 2:1–2:1, New York, NY, USA, 2011. ACM.
- [27] Ryan Poplin, Avinash V. Varadarajan, Katy Blumer, Yun Liu, Michael V. McConnell, Gregory S. Corrado, Lily Peng, and Dale R. Webster. Predicting cardiovascular risk factors from retinal fundus photographs using deep learning. *CoRR*, abs/1708.09843, 2017.
- [28] G. Preethi and V. Sornagopal. Mri image classification using glcm texture features. *Green Computing Communication and Electrical Engineering (ICGC-CEE)*, Mar 2014.
- [29] Lee Ann Remington. *Clinical Anatomy of the Visual System*. Elsevier Inc., St. Louis, MO, 2 edition, 2005.
- [30] L. Rokach and O. Maimon. Top-down induction of decision trees classifiers - a survey. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 35(4):476–487, Nov 2005.
- [31] R. Roslan and N. Jamil. Texture feature extraction using 2-d gabor filters. In *2012 International Symposium on Computer Applications and Industrial Electronics (ISCAIE)*, pages 173–178, Dec 2012.
- [32] C. W. Tan and A. Kumar. Accurate iris recognition at a distance using stabilized iris encoding and zernike moments phase features. *IEEE Transactions on Image Processing*, 23(9):3962–3974, Sept 2014.
- [33] Ramaiah University and S-VYASA. Iris diagnosis.
- [34] Gulshan V, Peng L, Coram M, and et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*, 316(22):2402–2410, 2016.
- [35] Wenchao Zhang, Shiguang Shan, Wen Gao, Xilin Chen, and Hongming Zhang. Local gabor binary pattern histogram sequence (lgbphs): a novel non-statistical model for face representation and recognition. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 1, pages 786–791 Vol. 1, Oct 2005.