Impact of Sensor Ageing on Finger-Image based Biometric Recognition Systems

Masterarbeit

Zur Erlangung des Diplomgrades an der
Naturwissenschaftlichen Fakultät
der Paris-Lodron-Universität Salzburg

Eingereicht von
Christof Kauba
ckauba@cosy.sbg.ac.at

Betreuer:
Univ.-Prof. Mag. Dr. Andreas Uhl

Department of Computer Sciences
University of Salzburg
Jakob–Haringer–Straße 2
5020 Salzburg
AUSTRIA

Salzburg, im Januar 2015

Abstract

Just like any other electronic device, image sensors show ageing effects over time. The most common effects are hot pixels and stuck pixels which are randomly appearing, permanent defects, increasing in number over time. These defects have the characteristics of spiky noise and degrade the image quality. The aim of this thesis is to investigate the impact of these ageing related defects on the performance of fingerprint, finger- and hand vein recognition systems. For this purpose, the recognition accuracy of different types of fingerprint (minutiae-based and non-minuatie based, e.g. correlation based), finger- and hand vein matchers (binarisation type, keypoint based and LBP based) is evaluated in terms of the EER. As it is not possible to distinguish the influence of sensor ageing related defects from other influences like noise, dirt, etc., the pixel defects are simulated using images containing no or only a negligible number of pixel defects as ground-truth. A pixel model is introduced and the images are aged using an ageing simulation algorithm. The simulation parameters (defect growth rate and defect parameters) are estimated using real sensor characteristics. The different fingerprint, finger vein and hand vein feature extraction and matching schemes are evaluated and compared against each other to find out if they are affected and how they are affected. This should give some clues about why they are affected and also what can be done to make them more robust against these ageing related defects. According to the experimental results, there is no considerable influence on either of the tested schemes for a realistic number of defective pixels.

Keywords: Image sensor ageing, pixel defects, defect detection, recognition accuracy, performance evaluation, EER, fingerprint recognition, finger vein recognition, hand vein recognition

Contents

1 Introduction

Nowadays each one of us has to authenticate oneself for many different reasons. E.g. if you want to withdraw money from the cash machine, you have to enter your personal PIN in combination with your ID card, the same if you switch on your mobile phone. At your home computer you might also need a password to log in and in other situations you authenticate via your signature. But all these traditional methods of authentication have some disadvantages. A signature can be forged, a PIN or password can be forgotten or even worse it can get disclosed and an ID card can get stolen. That is where biometric identification systems come into place. They all use some kind of biometric trait, which is unique, stable and cannot get lost or stolen. So there is no longer the need to remember complex passwords or carry ID cards. Biometric systems are therefore not only more secure but also more convenient for the users which lead to their widespread use, especially fingerprint recognition systems, because they are a mature technology and quite cheap. Usage scenarios include access control systems, screening at airports, fingerprint sensor at your home door and also the fingerprint sweep-scanners which are embedded into most modern notebook computers as an alternative to password log on.

If it comes to fingerprint recognition systems, there are many solutions available on the market suitable for different application scenarios. These are using different types of scanners, e.g. optical scanners, silicon based ones or thermal ones from which the optical scanners are the most prominent type. Also different kinds of feature extraction and matching schemes are used, most commonly employed are minutiae-based ones but especially for low quality fingerprint images non-conventional algorithms like correlation based ones are used.

An important factor regarding the recognition accuracy of a fingerprint identification system is the quality of the input fingerprint image. The quality can be degraded by the skin condition itself, e.g. dryness, moisture or dirt on the finger, by improper use, e.g. uneven pressure on the sensor or uneven sweeping motion but also by sensor related issues like noise and sensor ageing related pixel defects.

Moreover there are some scenarios where a fingerprint based recognition system cannot be used, e.g. time tracking for coal mine or construction workers as the dirt on their fingers and the abrasion of the fingers makes the use of fingerprint scanners practically impossible. A suitable alternative could be an iris based system, but the more practical and cheaper solutions are finger- and hand vein based systems. These systems have several advantages over fingerprint based ones. The veins are underneath the skin so the vein pattern is resistant against forgery as the veins are only visible in infrared light. Also liveness detection is easily possible. Moreover the vein patterns are neither sensitive to abrasion nor to finger surface conditions. But there are also some disadvantages. First of all, so far it is not completely clear whether vein patterns exhibit sufficiently distinctive features to reliably perform biometric identifications in large user groups. As the currently available data sets are limited in their size, this issue cannot be clearly answered at present state. Another major disadvantage is the capturing device which is, due to the required transillumination principle, rather big compared to a fingerprint sensor. Furthermore, the vein structure is influenced by temperature, physical activity and certain injuries and diseases. While the impact of these effects on vein recognition performance has not been investigated in detail so far, it is clear that suitable feature extraction methods should be independent of the vein width to compensate for corresponding variations.

There can also be other image distortions which influence the image quality. One of these are pixel defects related to ageing of the image sensor used as mentioned above. These defects are point-like and appear as spiky shot noise in the images. After thorough search of academic sources it has been determined that the impact of sensor ageing related pixel defects on the performance of fingerprint and finger- and hand vein based recognition systems, i.e. different feature extraction and matching schemes, has not yet been studied. Therefore, this evaluation is the topic of the present thesis.

1.1 Image Sensor Ageing

These days digital imagers are everywhere, starting from consumer cameras, mobile phone cameras, surveillance cameras. But also in biometrics digital imagers are often used to capture the biometric trait of a subject, e.g. fingerprints, finger- and hand vein images, iris images and face images.

The main part of a digital imager is its image sensor, which converts the incoming light into a digital signal, i.e. the output image. Just like any electronic device, an image sensor shows ageing related effects over time. An image sensor has a much bigger area compared to other electronic parts inside a digital imager, therefore it is more sensitive to external influences and especially to ageing processes. An image sensor consists of photosensitive cells, called pixels. Image sensor ageing leads to defective pixels, the main effects are so called hot pixels and stuck pixels. Although the probability for a single pixel to be defective is quite low, defective pixels occur in nearly every image sensor simply due to the high number of pixels an image sensor contains. These defective pixels are visible in the output image and if their number increases, the quality of the output image degrades. The effect is similar to spiky noise in the output image. Pixel defects are permanent, their number increases continuously and linearly with time and they are randomly distributed over the sensor area. A distinction is made between factory or fabrication time defects and in-field defects which are present immediately after the manufacturing of the image sensor or first appear while the sensor is in use, respectively.

These defects can be corrected by a factory calibration, at which the defective pixels are simply masked out, i.e. their output is not used, instead their value is replaced by an interpolation of the neighbouring good pixels. Factory mapping is not only expensive but also infeasible in many application scenarios, like embedded image sensors or extraterrestrial image sensors. Therefore, the impact of sensor ageing should be reduced to a minimum. The industry has not paid a lot of attention on the impact of sensor ageing related defects, because most consumer cameras are replaced after 3-4 years and during this time period the number of defects is rather low, except for factory time defects but these are always masked out in DLSRs and most consumer cameras. But the imaging devices used in biometrics are rarely replaced so the sensor ageing related defects may have a higher impact there.

1.2 Existing Literature Regarding Sensor Defects

Albert Theuwissen [3, 4] analysed the source causing the in-field sensor defects, mainly hot spots (hot pixels). He considered terrestrial cosmic ray radiation to be the main source of the pixel defects. So he did some experiments with sensors stored at ground level, others stored at elevated level and others were shipped on transatlantic flights. It is known that the intensity of the terrestrial cosmic rays is dependent on the altitude so if his assumption is true, there should be some differences in defect development of the 3 groups of sensors. Indeed he found differences which let him conclude that the neutrons of the terrestrial cosmic rays are the main source causing pixel defects in imaging sensors. Moreover, he performed some experiments at higher temperatures (storage and annealing experiments) and showed that storing the image sensors at higher temperatures than room temperature has a positive effect on the development of new defects (defects with high amplitudes did not occur any more). He also showed that it is possible to anneal some existing defects by storing the image sensor at temperatures of about 110°C for 24h. This reduces the number of hot spots independent of their amplitudes. The details of his experiments and results are described in section 3.3.

Chapman, Dudas, Leung et al. [5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17] did extensive research on several commercial cameras, i.e. DSLR, PS (point and shoot) and mobile phone cameras. They looked for sensor ageing related defects trying to establish cosmic ray radiation as the main causing source of defects. They searched for defective pixels not only manually by doing a dark-frame calibration on a regular basis and a manual inspection afterwards but also developed an automated defect tracing algorithm [14, 13] which is able to determine the date on which the defect first showed up using images taken by the sensor. This algorithm is explained in section 4.4. They found out that the main type of defects are hot pixels. Indeed they never found a true stuck pixel. Using the data collected during the measurements they did an analysis on the spatial and temporal distribution of the defects [16]. It became apparent that the defects are distributed uniformly over the sensor area and the size of the fault in the silicon lattice which causes the pixel to be defective is much smaller than the size of a single pixel (less than 2.3% of its size [13]), which indicates some type of constant stress, which the sensor is exposed to, rather than material degradation as defect source, else there would be defect clusters. It also turned out that the number of defects is increasing linearly with time (i.e. the defect rate is constant over time, implying a Poisson process) which also indicates some kind of causal mechanism like radiation as defect source rather than a single traumatic event or material degradation related defects, else the increase would be sudden or exponential, respectively. Thus they concluded that cosmic ray radiation might be the main source causing image sensor defects. In addition they also evaluated the influence of sensor parameters like sensor area, sensor type (CCD or CMOS), pixel size and ISO setting on the defect growth rate and found out that the growth rate scales linearly, which is another evidence for some kind of radiation as the source causing the defects. They also found out that reducing the pixel size has a great impact on the defect growth rate (smaller pixels lead to more defects) and that CCD sensors are more sensitive to the defect source than CMOS/APS ones, i.e. they develop more defects. Furthermore the ISO level has a big influence on the number of defects as the ISO setting is nothing else than a numerical gain factor, which also amplifies the defective pixel parameters, thus they become more visible at higher ISO levels and some of them might even be only detectable at a higher ISO setting. They were finally able to derive an empirical formula which is a power law relationship between the defect growth rate and the sensor characteristics. The details are described in section 3.7.

Jessica Fridrich [18] investigated imperfections of image sensors and how they can be used as a unique sensor fingerprint. Among these imperfections are manufacturing time defects, defects caused by physical processes inside the camera and also defects caused by environmental factors. She focused on the photo-response non-uniformity (PRNU) and the dark current. The PRNU is essentially a noise pattern caused by imperfections of the different pixels due to the manufacturing process (physical dimensions, quantum efficiency and homogeneity of the silicon may vary), which is overlaid on the image. Even if there is no exposure, a pixel may contain some free electrons causing a current flow due to thermal effects, commonly denoted as dark current. This dark current increases with rising temperature and higher exposure times. Therefore, the effect on the output image is dependent on the temperature, the exposure time but also the ISO setting. If the dark current of a pixel is very high, the pixel is called a hot pixel. If the pixel has a high additional offset, which is illumination independent, it is called stuck pixel. She proposed a pixel output model, which is described in section 3.4.1 and according to this model she derived the defect matrices based on multiple images taken with the same sensor using the maximum likelihood method. This is done by computing the noise residual using a custom-designed wavelet and a 3x3 median filter. She also found out that pixel defects occur randomly in time and space and are independent from each other. New defects appear with a constant rate, i.e. it is a Poisson process and once a defect occurs it is permanent. This is again an evidence of cosmic rays as the source causing defects.

1.3 Influence on Biometric Recognition Accuracy

Up until today, it is not clear if and how these sensor ageing effects influence the recognition accuracy of different biometric systems. To be able to quantify the impact on the recognition accuracy, at least two identical samples, captured at different points in time, of the same person showing no other ageing related effects than sensor ageing, would be needed. Unfortunately it is impossible to achieve identical conditions for both captures, as there are other influences, e.g. misalignment, human ageing and environmental conditions which make the separation of the sensor ageing related impact impossible. So it is not possible to capture real test data to evaluate the impact of sensor ageing.

This thesis proposes a method to evaluate the impact of sensor ageing on the recognition accuracy of finger image based biometric recognition systems, i.e. fingerprint, finger vein and hand vein based systems. To overcome the problem of other influences than sensor ageing related defects in the test data, an algorithm, which is able to simulate sensor ageing effects within images, is used to generate simulatively aged fingerprint, finger vein or hand vein images, respectively. Thus a common image data base can be used as a starting point for the experiments. As "ground-truth" the baseline matching performance, i.e. recognition accuracy, of the different matchers for this data base is determined first. This data base is then "aged" using the algorithm and the different matching schemes are evaluated using the aged versions of the images to measure the impact of sensor ageing on the biometric system's recognition accuracy. The parameters for the simulation (defect growth rate, types and intensities of the defects) are estimated using an empirical formula based on the biometric sensor's technical data. Then the results of the different matchers can be evaluated and the impact of the sensor ageing related defects can be quantified, compared against each other and also the differences can be investigated.

So at first, it reveals whether a single matcher is influenced by the ageing related defects or not, i.e. is the matching performance decreasing or is the matcher robust against it and the performance does not change at all. Through comparing to the results of the other matchers also the relation of these results can be evaluated to determine if all matchers are equally influenced or if there are some that are more robust against it than others. By doing a ranking of the matcher performances also the question of whether this ranking is influenced can be decided, i.e. is the best matcher still the best one under the influence of pixel defects or does the ranking change at some point? This leads to another important question, do all the matchers react in the same way to the defective pixels or are there any irregularities? So the main goal of this thesis is to investigate the influence of sensor ageing related pixel defects in fingerprint, finger- and hand vein images on the performance of different feature extraction and matching schemes. For fingerprint images 2 minutiae-based matchers, one correlation-based matcher and one ridge feature-based matcher have been evaluated. For finger- and hand vein images 5 binarisation-type approaches (trying to extract the vein pattern, creating a binary vein image and comparing the binary images then) and a keypoint-based approach (SIFT based) have been investigated.

The rest of this thesis is structured as follows: Section 2 describes the principle of an image sensor and also the two main sensor types, CCD and APS. Section 3 investigates the cause of sensor defects, describes the different defect types and gives two different pixel models describing these defects. In section 4 a defect detection algorithm, which is able to identify defects in conventional images (no dark-field or bright-field calibration images) is presented. Section 5 describes the simulation algorithm which is used to generate aged images by using given defect parameters. In section 6 the different fingerprint recognition methods are outlined. Section 7 describes the different feature extraction and matching methods used for finger- and hand vein recognition. Section 8 gives an overview of the experimental setup and the image databases. In section 9 the results of the different ageing experiments are presented and interpreted. Section 10 finally summarizes the findings of this thesis and gives an outlook on future work.

Below you will find an explanation of several terms and acronyms appearing in this thesis.

1.4 Acronyms and Terminology

APS Active Pixel Sensor, is one of the two major types of image sensors used in digital cameras and also biometric sensors.

CCD Charge-Coupled Device, is the second type of image sensor commonly used.

CMOS Complimentary Metal Oxide Semiconductor, is a technology for constructing integrated electronic circuits. It is also an abbreviation for the APS sensor which is sometimes called CMOS sensor.

DSLR Digital Single Lens Reflex camera, a digital camera combining the principle of a single-lens reflex camera with a digital image sensor.

ISO Level The ISO level or ISO setting originates from analogue photography and describes the sensitivity of a photographic film to the incoming light depending on the characteristics of the film. The ISO level of a digital image sensor is adjusted by setting the signal gain of the sensor. A higher ISO level means that the sensor is more sensitive to incoming light, i.e. images get brighter at the same exposure time. But with an higher ISO level also the noise level is increased.

MP Megapixel, a million pixels. Used to indicate the total number of pixels of an image but also for the maximum possible resolution of an image sensor.

PS Point-and-Shoot camera, the most common type of digital camera used in consumer area. It is designed for simple and easy to use operation. Due to its simple design compared to DSLR cameras these cameras are also cheaper than DSLRs.

Fingerprint, Finger- and Hand-Vein Images Digital images captured by a biometric sensor, i.e. a fingerprint scanner for fingerprint images, a finger vein scanner for finger vein images and a hand vein scanner for hand vein images, respectively. As the impact of sensor ageing related pixel defects should be investigated, only optical scanners are considered in this thesis, because the sensor ageing effects as described in section 3 are only applicable for optical image sensors. Fingerprint sensors are described in section 6.2.1. Finger- and hand vein scanners use an optical transillumination principle with which the finger or hand, respectively is screened using near-infrared light to render the veins visible. These scanners are explained in section 7.1. During the experiments regarding fingerprint matchers, the sample images from the Fingerprint Verification Contest 2004 (FVC2004, please refer to section 8.1.3 for details) were used. The experiments on finger vein recognition were conducted on images of the University of Twente Finger Vascular Pattern Database (UTFVP, details in section 8.2). For the evaluation of hand vein recognition performance a data set captured at the University of Salzburg (see section 8.3) was utilized.

Recognition Performance The different matchers were evaluated using the sample data mentioned before. The recognition performance was quantified in terms of the matching performance. Therefore, the procedure of the FVC2004 was adopted, which is described in section 8. In order to be able to compute the matching performance, the matching results or matching scores, respectively, were analysed and the following numbers were calculated:

- False Match Rate (FMR)
- False Non Match Rate (FNMR)
- Equal Error Rate (EER)

Biometric Identification A biometric identification system tries to determine the identity of a subject. It is given a specific kind of biometric trait, e.g. a fingerprint, and compares this with all the samples previously enrolled and stored in the data base. It returns the best match including the matching score and should also return the likelihood that the found match is correct, depending on the impostor matching scores and the sole genuine one. Such a system performs an 1 to n match, i.e. it matches the given input to all other samples stored in the data base.

Biometric Verification Contrasting to identification systems, a verification system is not only given a biometric trait but also information about the claimed identity. It thus only has to compare the trait with the one stored in the data base under the same ID. If the matching score is above a preset threshold, the system confirms the claimed identity, otherwise not. The first one is called a match, the second case is called non-match. A verification system performs an 1 to 1 match, i.e. it matches the given input only to the sample stored in the database with the given ID, at least in theory.

Matching Score A biometric recognition system compares two samples of the same biometric trait, e.g. it gets two fingerprint images as input and the output of the system is a matching score. This score indicates either the similarity or the difference between the two inputs. Depending on this score the result of an identification or a verification process, respectively is determined. Using a threshold t and comparing the matching score s to this threshold, the output can either be a match ($s \geq t$) or a non-match ($s < t$, vice versa if the score value indicates difference instead of similarity). Usually the matching scores are normalised to be in a range of $[0, 1]$.

Genuine Score A genuine score is a resulting matching score of a comparison in which the given input corresponds to the sample stored in the data base with the same identity as the claimed input one.

Impostor Score An impostor score is a resulting matching score of a comparison in which the given input does not correspond to the sample stored in the data base with claimed input identity.

False Match Rate (FMR) A false match occurs when two samples taken from a different subject are incorrectly classified as originating from the same subject, e.g. two fingerprints from different persons or fingers are declared to be from the same finger. Given a whole data set, the false match rate denotes the number of false matches divided by the total number of performed matches.

False Non Match Rate (FNMR) A false non match occurs when two samples are classified to be from two different subjects but actually they belong to the same subject, e.g. two fingerprints taken from the same finger are declared to be from different fingers. The false non match rate denotes the number of false non matches divided by the total number of performed matches for a given data set.

Equal Error Rate (EER) By adjusting the threshold within a set of matching scores, the matching scores will either lead to a match or a non-match, both the FMR and FNMR change. The EER is the point where the FMR and FNMR have the same value. It is a performance indicator for a biometric recognition system. A low EER corresponds to low FMR and FNMR values and thus to a high recognition accuracy.

Simulation of Sensor Ageing During this work only the influence of sensor ageing related pixel defects and no other external influences should be evaluated. Real fingerprint finger vein or hand vein images will always contain other distortions due to finger surface conditions, dirt on the sensor area, subject ageing, etc. Thus, in order to be able to exclude these external influences the sensor ageing effects are simulated, i.e. the pixel defects are generated artificially and overlaid onto the images using the algorithm described in section 5. Therefore, the images of the sample data sets mentioned above were used as a basis, i.e. as unaged images and the algorithm was applied to generate several data sets with differently aged versions of the images, which were then used to determine the matching performance.

Test-Run A test-run describes a run of all the matches necessary for determining the EER (genuine and impostor matches) according to the test protocol adopted from the FVC2004 for a single matcher and a single aged data set. For the fingerprint data this results in a total of 7750 matches performed. For the finger vein data 12740 matches and for the hand vein data set 5250 ones are performed. Afterwards, the resulting EER based on the match scores is calculated which is the result of the test-run.

2 Image Sensors

For an accurate analysis of the impact of sensor defects, at first some of the fundamental principles of solid-state image sensors should be explained. This section gives a short overview of the basic working principle of an image sensor from light detection to digital cameras. In short an image sensor consists of an array of photosensitive cells, called pixels, and converts the incoming light first into a proportional voltage or current which is then converted into a digital value using an analog to digital converter. Although the rest of the imager system is digital, the sensor is an analog device. In an ideal image sensor every pixel would be rectangular and have a unified photo-response, which means each pixel produces the same output at the same level of incoming light. But in reality there are some pixels that are more sensitive compared to others due to imperfections in the manufacturing process, these are called manufacturing time defects. Moreover there are also other effects, like the temperature dependent dark current which affect the pixels' output. Even worse some of the pixels may change their photo-response characteristics over time, which are typically called in-field defects. This effect is commonly called sensor ageing.

Today there are two main types of image sensors, CCD and CMOS which are different in the way how they convert the photons from the incoming light into electrical signals. Each of these two has its advantages but also disadvantages which are described below.

2.1 Photodiodes and Photogates

The light detection mechanism is shared by both CCD and CMOS sensors. At the very low level the photoelectric effect is the basis of light detection. If the energy of a photon from the incoming light is large enough, its energy is absorbed by the semiconductor material, which leads to the elevation of an electron into the conduction band, generating a free carrier. On its own this is not sufficient to convert the incoming light into an electric current as the free carriers recombine after a short time.

The freed electrons have to be captured in an electric field such as it is used in a photodiode. A photodiode is a semiconductor device, much like a regular semiconductor diode, containing a p-n junction, which converts the incident light into an electric current. By applying a positive voltage it is operated in the so called photo-conductive mode, generating a depletion layer, which also reduces the response time. The voltage is then removed and light which reaches through to the depletion region generates free electron-hole pairs in it. Due to the built-in electric field these pairs are separated and the holes and electrons are pushed to the opposite ends of the photodiode which induces an electric current. As this current is very small, this process lasts only for a short time, which is called the integration time. During this time the charge is collected inside the photodiode, which is then read, followed by a reset phase at which the positive voltage is applied again. Ideally the generated photocurrent should be dependent on the incident illumination only, but even in an ideal photodiode there is a so called dark current. The dark current is a thermally generated leakage current due to the applied reverse bias voltage, which is dependent on the temperature and the junction width, leading to a discharge even in complete darkness. Due to impurities occurring during the production process, causing defects in the silicon lattice, this thermally generated current is not uniform and may be generated more rapidly.

A photogate is used to convert light into an electric current and is a metal oxide semiconductor capacitor. Its operation is similar to a photodiode, only the construction is different. It consists of a p-substrate, an oxide layer and a gate electrode above this layer. Again to create a depletion region underneath the gate, a positive voltage is applied to the electrode. Inside this region the photons contained in the incident light create the free electron-hole pairs during the integration time and the current is read out at the end of the integration time.

2.2 CCD

During early development CCD was the most prominent type of image sensors. A CCD sensor basically consists of a series of MOS capacitors that convert the incoming photons (illumination) into electron charges at the semiconductor-oxide interface. These charges are then shifted sequentially between the capacitive bins inside the device for readout. Due to this shift, which is necessary for readout, CCD sensors are limited in pixel transfer speed and thus in frame rate. An advantage of CCD sensors is their good optical response even in low light conditions.

2.3 CMOS

CMOS sensors are also called APS (active pixel sensors) because in contrast to the CCD sensors each pixel contains of a photodetector (photodiode or photogate) and an active amplifier, i.e. a transistor. The additional amplifier reduces the fill factor ($\sim 25-30\%$) but makes faster addressing and readout possible, thus allowing higher frame speeds. Nowadays so called micro lenses are used which cover the whole area of a pixel and focus the incoming light onto the photoactive area, therefore the efficiency of conversion is nearly the same as with CCD sensors. Another disadvantage is that since a CMOS sensor captures one row at a time there may be a "rolling shutter" effect, which skews the image. Another issue is amplifier variation. Each pixel has a built-in amplifier, these amplifiers might not all have the same gain factor and therefore the output of the pixels is different even at an uniform input.

The biggest advantage of APS sensors is that they are CMOS compatible, therefore they can be seamlessly integrated into embedded devices and are also able to combine the sensor function with integrated image processing functions. Another advantage is that they are less expensive compared to CCD sensors. Moreover, CMOS sensors have a lower power consumption.

2.4 Digital Cameras

A digital camera typically consists of an optical system, including a lens and an optics board, an image processing board containing the image sensor, a housing, a LCD display and some kind of storage device. There are several types of digital cameras, the most prominent ones are the DSLRs (Digital Single Lens Reflex) and the PS (point and shoot) cameras.

2.4.1 CFA - Colour Filter Array

Each pixel is sensitive to light in a specific wavelength band, typically from 400 to 700 nm if the photosensitive cell is made from silicon. Pixels are monochromatic, i.e. they do not generate colour information, they just convert the incoming light into an electrical representation. As images are captured in RGB format, a way to make the pixels only sensitive for incoming light in the red, green and blue wavelength band, respectively has to be found. The most common way to achieve this is the use of a colour filter array (CFA). An array of colour sensitive filters is placed on top of the sensor in a way that at every pixel only red, green or blue light is able to pass through the filter. The most widely used filter is the so called Bayer pattern, which uses 2 green, 1 red and 1 blue filter, as shown in figure 1.

The problem which now arises is that every pixel now only captures information in one colour band. To get a whole colour image, i.e. red, green and blue colour information at every pixel, the missing information of the other colour bands at that pixel has to be interpolated. This colour interpolation is also known as demosaicing. Current demosaicing algorithms do not handle defective pixels, so the impact of the ageing related defects on colour images is higher due to demosaicing. Kimmel type demosaicing algorithms (adaptive) are most commonly used in commercial cameras.

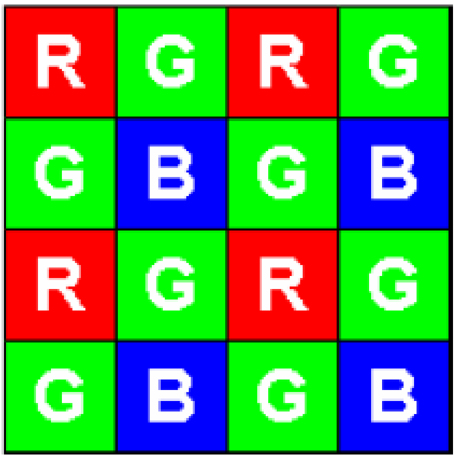

Fig. 1: Bayer CFA pattern [14]

Another problem with the colour filter array is the reduction of the overall sensitivity of the sensor. As with each filter only one colour band is allowed to pass through and reach the sensor surface, the remaining components of the incoming illumination are lost. Only 50% of the green and 25% of the blue and red intensity is retained.

3 Sensor Defects

Digital image sensors enjoy a widespread use today and become more and more popular. Starting from consumer cameras to mobile phones, cars and also other consumer electronic devices such as notebook computers, TV sets, etc. Like every creature also electronic components age over time. Unlike other electronic components, image sensors show ageing related defects soon after fabrication. Sensor defects occur in the form of defective pixels, which show a different characteristic than at manufacturing time. Some are completely insensitive to incoming light, others add an offset and are more sensitive to illumination than the other pixels. The accumulation of these defective pixels degrades the quality of the output image but fortunately they do not prevent the sensor from still providing useful output data. The main source causing these defective pixels is radiation. Unlike other devices, at which radiation causes mainly transient soft errors, i.e. errors that recover after a sensor reset or a certain amount of time, these image sensor ageing related pixel defects are permanent, i.e. they do not heal over time, their characteristics do not change and their number increases continuously with time.

In this chapter the different types of pixel defects are explained and a pixel model to describe the defective pixels is presented. The spatial and temporal distribution of the defects is studied, which gives some hints indicating the source causing the defects. The effect of colour demosaicing on single pixel faults is analysed, followed by an overview of different defect identification techniques. Finally the impact of sensor parameters on the defect growth rate is investigated.

3.1 Material Degradation Related Defects

Defects related to material degradation are mainly caused by to limitations and impurities during the manufacturing process of the sensor. These defects make the sensor more vulnerable and lead to early material breakdown. There can be minor faults like errors in the output image but also severe faults like a complete failure of the sensor.

Manufacture time defects are those which occur before the image sensor leaves the factory. Such defects exist but as the manufacturer does a so called factory calibration, i.e. generating a defect map and using this map to mask out the defective pixels during the colour interpolation (demosaicing) step and additionally by performing a dark frame subtraction after each image capture, these defects are not noticeable.

Defects related to material degradation that occur in-field are e.g. gate oxide thinning, hot carriers and electromigration. These defects are localized in a small area and will usually result in a cluster of spatially close defects which all develop at around the same time. The number of material degradation related defects rises exponentially with time and as it is shown in section 3.3.1 and 3.3.2 that both conditions do not apply for typical in-field pixel defects, material degradation cannot be the main source causing in-field defects. Therefore these defects are not explained in more detail here.

3.2 In-Field Defects

In contrast to manufacture time defects, in-field defects develop while the sensor is in use. But they not only develop while it is in use, also while it is stored. The most prominent types of in-field defects are single pixel failures like hot or stuck pixels. The impact of these defects could be overcome by performing a factory calibration. This is not only expensive but also infeasible for imagers in embedded devices and in remote sensing applications. Some researchers [13] found exclusively hot pixels and no other types of defects, which indicates that the source causing hot pixel defects does not lead to other defect types like stuck and abnormal sensitivity defects. Details of the defect types are described below.

There are other types of in-field defects related to electrical stress, e.g. fluctuations in the supply voltage and also electrostatic discharge, which are not discussed here.

3.3 Mechanism Causing the Defects

According to the literature, the main mechanism causing sensor defects is cosmic ray radiation.

Albert Theuwissen [3, 4] studied the influence of terrestrial cosmic rays on the number and parameters of newly generated pixel defects in image sensors. Terrestrial cosmic rays are the result of high energy particles originating from space (primary cosmic rays) which hit the atmosphere and produce secondary particles (secondary cosmic rays). He showed that the main mechanism causing the defects are neutrons, which are part of the terrestrial cosmic rays, that lead to displacement damage in the silicon bulk. The energy and density of cosmic rays is dependent on the altitude, the latitude as well as on the earth's magnetic field. During an air plane flight, the density of neutrons in the cosmic ray total flux is about 100-300 times higher compared to ground level. To show that the terrestrial cosmic rays are the cause of the sensor defects, he evaluated several image sensors which were stored on-the-shelf, others were kept running and others were shipped around the world in air planes. As the radiation due to cosmic rays is higher during transatlantic flights, sensors shipped by air plane should show a higher defect rate. He indeed showed that there is an increase in the hot spot density of the shipped sensors compared to the stored ones (about 100 times higher which is in correlation with the increase in neutron density). Moreover he discovered that the probability to develop hot pixels with high amplitudes is relatively higher for sensors after the flights. He also stored some sensors at elevated altitudes and discovered what was expected, an increase in the number of hot spots compared to the sensors stored at sea level. In addition he stored some sensors underground, which should lead to an improvement, i.e. a lower number of hot spots, but the results of his experiment did not confirm this, which indicates that cosmic rays are not the only source generating new pixel defects. Leung et al. [13, 14] also noticed a similar increase in defect growth rate or defect count, respectively, after transatlantic flights present at their tested cameras.

Another mechanism causing sensor defects might be the increase of the dark current level over time. If the dark current increases, also the amplitudes of the hot pixels increase and therefore they become more visible.

Albert Theuwissen [3] concluded from his experiments that hot pixels with high amplitudes are mainly created due to the damage in the bulk of the silicon substrate caused by terrestrial cosmic rays. A high-energy neutron of the cosmic rays may hit a silicon atom right in its centre, which leads to a displacement of this atom from its rigid crystal structure, creating an interstitial and a vacancy, which are unstable. They migrate to energetically favourable positions in the lattice and may become trapped near impurity atoms due to the stress imposed on the lattice by the impurities and remain in a stable position there, creating a hot spot [4]. Hot pixels with lower amplitudes are most likely caused due to damage created at the Si-SiO2 interface, which results in an overall increase of the dark current. He also showed that the creation of new hot pixels is independent of technology, architecture, sensor type and vendor, but the amplitude of the hot pixels can depend on the sensor parameters.

In the second part of his experiments [4] he showed that also the storage temperature has an influence on the generation of new pixel defects. According to his results storing an image sensor at higher temperatures than room temperature results in an overall improvement, i.e. a reduction of the number of hot pixels with a high amplitude and reduces the generation of defects considerably. At 180° C hot pixels with a high amplitude no longer show up but the number of hot pixels with a small amplitude increases. The best storage temperature for image sensors is between 60° C and 110° C. He concludes that there has to be an instant annealing effect of hot pixels generated by cosmic rays but on the other hand another mechanism causing the generation of new hot pixels at higher temperatures.

In addition he performed annealing experiments with sensors usually stored at room temperature and found out that annealing at 100° C for 24 h seems to be very efficient in reducing the number of hotspots. This statement is true for hot spots of all amplitudes [4]. Therefore an anneal of 24 h at relatively low temperatures can be a good alternative to the storage at elevated temperatures. The overall effect seems to be the same. After 24h at 110° C most of the hot spots are annealed [4]. So this means that the hot pixels are not completely permanent as they can be annealed at higher temperatures. This and the influence of higher temperatures can by explained by the fact that the mobility of the vacancies and interstitial silicon atoms increases with increasing temperature, so the chance that a vacancy and interstitial recombine becomes higher. Moreover, the energy of the vacancies becomes larger so that a vacancy is instantly released from the trap again.

3.3.1 Spatial Distribution of Defects

The spatial distribution of the defects across the sensor area follows a normal random distribution (also true among different ISO levels) which means that the defect source is neither related to the manufacturing process, e.g. the design of the pixels, nor to material degradation. Material degradation would lead to defect clusters as described above.

Leung et al. [13, 16] applied statistical analysis to the defect distribution, showing that it is indeed a random distribution. First they used the defect map and inter-defect distance distribution and compared it with a normal distribution. They showed that there was no significant bias towards short or long distances [7], which would be the case for defect clustering. This can be seen in the histograms for defect distances of APS and CCD image sensors, shown in figures 3 and 4. This was followed by a chi-squared goodness of fit test, which also showed that the hypothesis was correct. In addition they did a distance to the nearest defective neighbour analysis with subsequent statistical tests (Z-score) and finally a Monte-Carlo simulation by which they found out that the minimum distance between two defects is 17 µm (at a pixel size of 6-7 µm) [14], inter-defect distances follow a broad distribution and the average distance is 10 mm [13]. This again supports the hypothesis that the defects are indeed randomly distributed. In other words no single event and also no material degradation related effect can be the source causing the defects. They [13] also estimated the size of the defect creating a hot pixel using statistical methods and found out that it is very small ($<$0.07 µm [19]/0.2 µm [13]/$<$0.04 µm [20]) compared to a usual pixel size of 2.2 µm, i.e. the defects are nearly point like, which contributes to the hypothesis that the defects are isolated and not clustered and therefore not caused by material degradation. As the defects are caused by a point like source, the overall dark current magnitude of the pixel should remain the same independent of the pixel size.

Albert Theuwissen [3] showed that the number of defects is increasing if the camera is exposed to higher cosmic ray radiation (e.g. transatlantic flights), which is another indicator that the main source of the in-field defects is cosmic ray radiation. He also found out that the neutrons of the cosmic ray radiation are causing these defects.

3.3.2 Temporal Distribution of Defects

Leung et al. [14] analysed the temporal distribution of the defects using two different methods. First they did dark-frame calibrations at regular time intervals (i.e. on a yearly basis) followed by a manual calibration and second they used an automated defect trace algorithm to determine the point in time when the defect first occurred. This algorithm works with regular scene images and is described in section 4.4.

If the defects are caused by a causal mechanism, i.e. constant stress to the image sensor, like cosmic ray radiation, the number of defects should increase linearly with time, contrary to a sudden increase if they were caused by a single traumatic event, like a shock or an exponential increase, i.e.

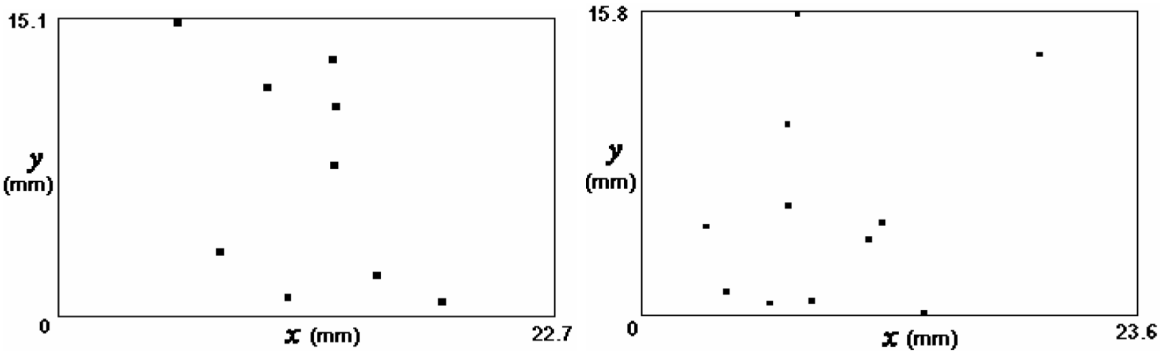

Fig. 2: Defect map APS (left) and CCD (right) [13]

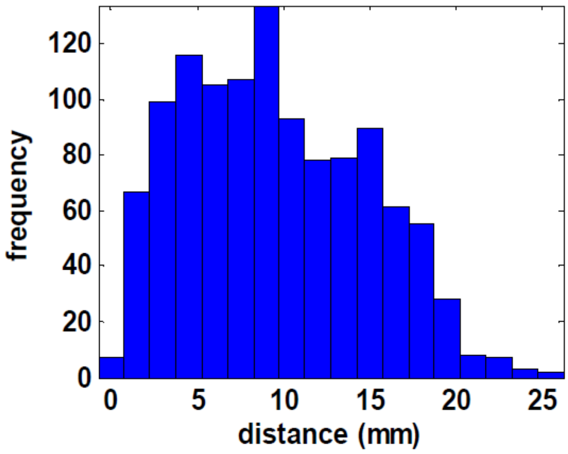

Fig. 3: Defect distance distribution for APS and CCD imagers [16]

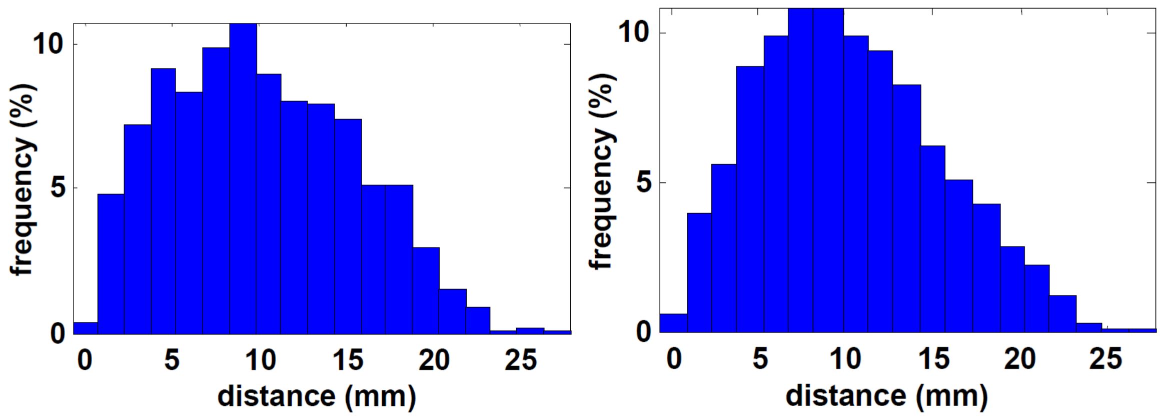

Fig. 4: Defect distance distributions for ISO400 (left) and ISO1600 (right) [17]

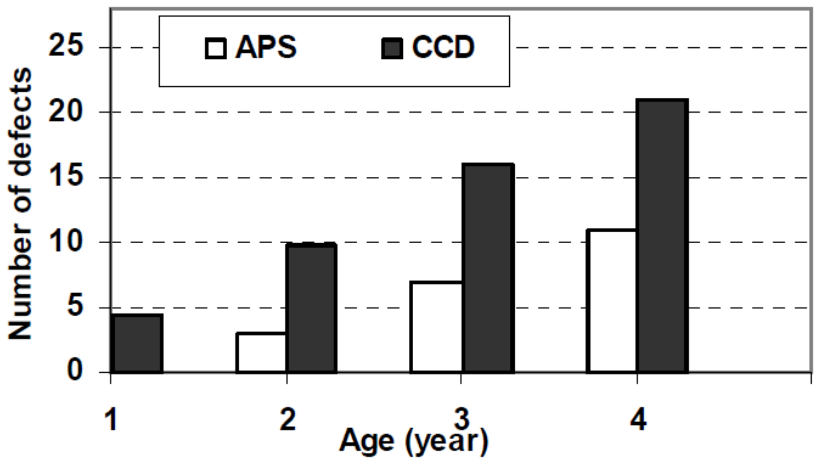

Fig. 5: Temporal growth of defects [16]

the time between two consecutive defects gets shorter with increasing age, if they were caused by material degradation, respectively.

They showed that the defects were indeed increasing in number linearly with time, that the defects are permanent, which means if a defect is there it will not disappear and the defect parameters do not change over time. They also analysed the inter-defect times, i.e. the timespan between two consecutive occurring pixel defects, using statistical methods, which lead to the result that the inter-defect times are following an exponential distribution. This indicates a constant defect rate which can be modelled by a Poisson process and is in contradiction to material degradation as defect source because then the defect rate would increase with time. Another important result is that the sensitivity to defects does not increase with time [7]. Moreover they calculated the defect growth rate for several types of imagers and Chapman et al. [9] carried on their work by developing a formula showing the influence of the imager properties on the growth rate, which is explained in section 3.7.

3.4 Defect Types

According to the more recent literature there are several defect types which can be distinguished by their photoresponse. Table 1 taken from [21] summarizes the most prominent defect types which occur either in-field or during the manufacturing process.

3.4.1 Pixel Defect Model (Fridrich)

A pixel model considers the incoming illumination and the impact of pixel defects on the raw output of the pixel or the sensor, respectively. Jessica Fridrich [18] used the following pixel model:

$$Y = I + I \circ K + \tau D + C + \Theta$$

$$with\ \ Y, I, K, D, C, \Theta \in \mathbb{R}^{w \times h}; \tau \in \mathbb{R}\ \ and\ w, h \in \mathbb{Z}$$

where Y is the sensor output, i.e. the image, I is the intensity of the incoming light (incident illumination), $I \circ K$ is the photo-response non-uniformity PRNU, τD the dark current (with τ being a multiplicative factor taking into account the exposure setting, sensor temperature, ...), C is a light-independent offset and Θ is some additive modelling noise.

Responsive to light	Defect type	Output function	Description
NO	Stuck high	$f(x) = 1$	Appears as a bright pixel at all the time
	Stuck low	$f(x) = 0$	Appears as a dark pixel at all the time
YES	Partially-stuck	$f(x) = x + b$	Offset $0 < b < 1$
	Hot pixel (standard)	$f(x) = x + R_{dark} \cdot T_{exp}$	Illumination independent offset that increases linearly with exposure time, i.e. I_{dark}
	Hot pixel (partially-stuck)	$f(x) = x + R_{dark} \cdot T_{exp} + b$	Has two illumination independet offsets: (1) increases with exposure time, R_{dark} (2) offset at all the time, b

Tab. 1: Characteristics of different defect types

Fig. 6: Defective pixel example (left: whole image, right: detail of the defect region) [18]

According to this model, a pixel with an extremely high dark current value D is called a hot pixel and is the most common point defect occurring on an image sensor. Another defect type, the stuck pixel, has a high offset value C. Both pixel defects occur randomly, are uniformly distributed over the sensor area and are independent from each other. Coming to sensor defects the interesting thing are the two defect matrices D and C, which can be estimated using several images taken by the same sensor. This is done by applying the maximum likelihood method and by working with the noise residuals $W = Y - F(Y) = I \circ K + \tau D + C + I - F(Y) + \Theta$, obtained using a denoising filter F, to improve the signal-to-noise ratio. As hot and stuck pixels are spiky in nature, a non-linear filter, like a median filter is a good choice for extracting the spiky pattern correctly.

$Y_k,\ k = 1...d$ are d images taken from regular scenes at known points in time. Using the noise residuals and the pixel model above:

$$W_k(i) = I_k(i) \circ K(i) + \tau_k D(i) + c(i) + \Xi_k(i)$$

with $\Xi_k(i) = I_k(i) - F(Y_k(i)) + \Theta_k(i)$, modelled as an i.i.d. Gaussian sequence with zero mean and variance $\sigma^2(i)$, where k and i are the image and pixel indices, respectively. According to the maximum likelihood principle, the unknown parameter vector $\Theta = (K, D, c, \sigma)$ for a fixed pixel i can be estimated using (for details see [18]):

$$\hat{\theta} = \underset{\theta}{arg\,max}\, L(W_1, ..., W_d|\theta)$$

Simplification of this Model by Bergmüller et. al For this work I adopted the pixel model of Bergmüller et al. [22] which is a simplified version of Jessica Fridrich's pixel model and show that it is similar to the one proposed by Dudas et al. Even if the pixel models of Dudas et al. and Jessica Fridrich seem quite different at first sight, they are not. Dudas et al. simply do not include the PRNU and the additional modelling noise.

Since all pixels are independent and all operations are done element-wise, the matrix elements $y_{x,y} \in Y$ are denoted as $y \in Y$ for simplicity, the same for $i \in I$, $k \in K$, $d \in D$, $c \in C$ and $\theta \in \Theta$. All age independent effects can be eliminated since I am interested in the ageing effect of one specific sensor, so the PRNU can be eliminated. As I aim for reproducible tests, modelling noise and environmental influences should be minimized, in fact they can be eliminated completely in a simulation, therefore $k = \theta = 0$. For all images taken with the sensor the same exposure settings are used (typically true for fingerprint, finger vein and hand vein scanners), therefore $\tau = const.$ and $\tau = 1$ for simplicity. According to the literature the dark current level is very low for short exposure times which are normally used for standard photographs but also for fingerprint images to avoid motion blur. Taking all this into account a simplified pixel model can be derived:

$$y = i + d + c \;\; with \;\; y, i, c, d \in \mathbb{R}$$

The most prominent defect types that develop over a sensor's lifetime are hot and stuck pixels. If the dark current d of a pixel is extremely high it is often denoted as hot pixel, whereas if the offset c is high this results in a saturated pixel and is denoted as a stuck pixel then. As the definitions in the literature are not consistent, Bergmüller et al. [22] defined the following model for pixel defects:

$$y = c$$

$$y = i + d$$

where the first one is light independent and has a constant value c, denoted as stuck pixel and the second one adds an offset to the incident illumination and is referred either as partially-stuck or hot pixel. The cause for hot pixels is a higher dark current at that pixel compared to others. The dark current level depends on the temperature and exposure time, which are both kept constant in the experiments, thus the dark current level is constant and therefore there is no difference between a hot and a partially-stuck pixel, thus it is simply denoted as hot pixel. This model for stuck pixels can be directly compared to the one used by Dudas et al. If one takes their hot pixel model (see Equation 1), set $m = 1$, $T_{Exp} = 1$, $R_{Dark} = 0$ and set $I_{Pixel} = y$, $R_{Photo} = i$ and $b = d$ (as discussed before), this reveals the same hot pixel model as it was derived above.

This leads us to the following pixel model for 8 bit grey-scale images:

$$Y(x,y) = \begin{cases} C(x,y) & if\, C(x,y) \neq 0 \\ I(x,y) + D(x,y) & otherwise \end{cases}$$

$$with\;\; Y, C, I, D \in (\mathbb{Z} : [0; 255]^{w \times h})$$

where C and D are the defect matrices. A pixel's output $Y(x,y)$ saturates at 0 and 255 if interval borders are exceeded.

This pixel model is the basis for the ageing simulation algorithm, described in section 5.

3.4.2 Pixel Defect Model (Dudas and Leung)

During their first experiments, Dudas et al. [12] modelled several defect types using the following equations, where the output range of a pixel is $0-1$ and $x = I_{photo} \cdot T_{integration}$ measures the incident illumination:

$$f_{Good}(x) = x$$

$$f_{Stuck-Low}(x) = 0$$

$$f_{Stuck-High}(x) = 1$$

$$f_{Stuck-Mid}(x) = c, \;\; 0 < c < 1$$

$$f_{Partially-Stuck}(x) = x + b$$

$$f_{Abnormal-Sensitivity}(x) = m \cdot x + b$$

$$f_{Hot-Pixel}(x) = m \cdot x + b + I_{Hot} T_{Integration}$$

But afterwards they simplified their model to only include hot pixels and partially-stuck hot pixels. They [12, 17] modelled the response I of a pixel using the following equation, which also includes the ISO level:

$$I_{Pixel}(R_{photo}, R_{Dark}, T_{exp}, b) = m \cdot (R_{photo} \cdot T_{exp} + R_{Dark} \cdot T_{exp} + b) \tag{1}$$

where R_{photo} measures the incident illumination, R_{Dark} is the dark current rate, T_{exp} is the exposure time, b is the dark offset and m is the amplification proportional to the ISO setting.

	Fridrich [18]	Bergmüller [22]	Dudas [12, 17]
Model	$Y = I + I \circ K + \tau D + C + \Theta$	$y = i + d + c$	$I_{Pixel}(R_{photo}, R_{Dark}, T_{exp}, b) =$ $m \cdot (R_{photo} \cdot T_{exp} + R_{Dark} \cdot T_{exp} + b)$ $f_{Stuck-Mid}(x) = c,\ \ 0 < c < 1$
Sensor/Pixel	Whole sensor	Whole sensor	Single pixel
Sensor/Pixel output	Y	y	I_{Pixel}
Noise	Θ	n.a.	n.a.
PRNU	$I \circ K$	n.a.	n.a.
Dark current	τD	d	$R_{Dark} \cdot T_{exp}$
Illumination	I	i	$R_{photo} \cdot T_{exp}$
Hot pixel	$D >>$	$y = i + d$	$I_{Pixel} = m \cdot (R_{photo} \cdot T_{exp} + R_{Dark} \cdot T_{exp})$
Stuck pixel	$C >>$	$y = c$	$f_{Stuck} = c,\ \ 0 < c < 1$

Tab. 2: Comparison of pixel defect models

For an ideal good pixel, both R_{Dark} and b are 0 and the output is only proportional to the incident illumination. A hot pixel now adds a signal on top of the pixel's output, therefore the output of a defective pixel will appear brighter.

As it can be seen, this pixel model only includes hot pixels and partially-stuck hot pixels. Dudas et al. did measurements over the last 7 years and they never found a true stuck pixel, even though they are discussed in the literature. Moreover they did not find any abnormal sensitivity pixels. Instead they found partially-stuck hot pixels with a high offset, appearing as stuck high pixels. Therefore they use this simple pixel model, which only considers hot and partially-stuck hot pixels. During their measurements [5] they also found out that nearly 80% of the hot pixels are of the partially-stuck type.

3.4.3 Comparison of Defect Models

Table 2 compares the defect models of Jessica Fridrich, Bergmüller et al. and Dudas et al. Please note that some of the entities listed in the table cannot be directly compared, i.e. their units do not match, but they are affecting the output of the sensor or a single pixel in the same way.

3.4.4 Stuck-High

A stuck high defect is a pixel which output always has the same fixed output value, independent from the incoming illumination. A stuck high pixel will appear always bright in the output, i.e. it is stuck on a high value, according to the pixel model above this means the value 1. Dudas and Leung et al. [12, 17] never found a true stuck-high pixel. Instead they found partially-stuck hot pixels with a high offset (especially if a high ISO settings is chosen) so they suggest that the development of stuck high pixels in the field may actually be due to the presence of hot pixels with high offsets.

3.4.5 Stuck-Low

A stuck low defect is basically the same as a stuck high one, except that its fixed output value is 0, which means this pixel always appears as a dark spot dark in the output image. If cosmic ray radiation is the main source causing the defects and if the mechanisms are like Albert Theuwissen [3, 4, 23] described, then a stuck-low pixel cannot be explained by cosmic ray radiation as source causing the defects. Defective pixels due to radiation damage can only appear brighter in the output image.

According to Dudas et al. [12], stuck-high and stuck-low pixels are solely factory time defects, which are corrected using factory time mapping. Therefore, they have not found a single stuck defect

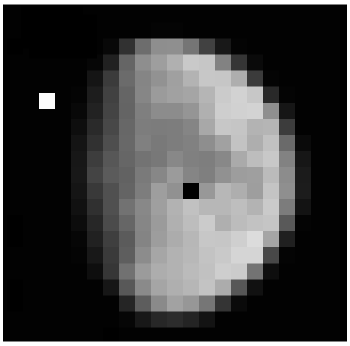

Fig. 7: Stuck low and stuck high pixel [11]

during their tests.

3.4.6 Stuck-Mid or Fully-Stuck

A stuck-mid [12] or also called fully-stuck [13] pixel is a generalised form of stuck-high and stuck-low pixels, which is stuck at an arbitrary but fixed value c in the range $0 \leq c \leq 1$. Thus a fully-stuck pixel will always have the same output under all illuminations. As with the stuck-low and stuck-high pixels, Dudas and Leung at al. never found a true fully-stuck pixel. A fully-stuck pixel is therefore also assumed to be a factory time defect, which is corrected by factory mapping.

3.4.7 Partially-Stuck

A partially-stuck pixel has an illumination independent offset b like a fully stuck one, but also an illumination dependent component. Therefore the output of a partially-stuck pixel is not always the same under all illuminations. It depends on the exposure time. Again Dudas et al. [12] found out that all partially-stuck defects that they described earlier were actually also hot pixels, which suggests that a common mechanism may lead to both defects.

3.4.8 Abnormal Sensitivity

There are basically two types of abnormal sensitivity pixels, low sensitivity ones, showing only a fraction of their incident illumination and high sensitivity ones, showing more than the corresponding incident illumination on the output. This type of defect is identified by comparing the illumination response of the pixel with its neighbours. As the neighbouring pixels should have the same sensitivity, variations in sensitivity indicate the presence of an abnormal sensitivity pixel defect. Once more Dudas et al. [13, 12] never found an abnormal sensitivity defect.

3.4.9 Hot Pixel

A hot pixel is a defect that has an illumination independent component which increases linearly with exposure time. Whereas the dark response of a good pixel should be close to 0 (due to sensor noise it is not exactly 0), the dark response of a hot pixel increases with the exposure time. The output of a hot pixel is different from noise and it appears as a bright spot with a fixed location in the output image. The added dark response limits the light collection capacity of the pixel, causing it to saturate at a lower illumination level, which reduces the pixel's dynamic range. As a hot pixel results from non-light generated charge at the pixel, i.e. an increased dark current, it is also called gain enhanced

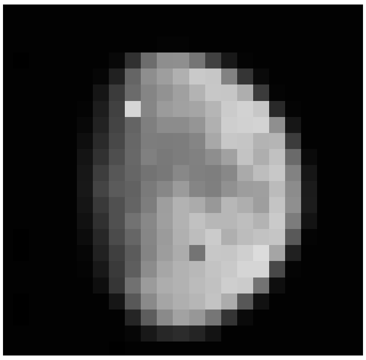

Fig. 8: Abnormal sensitivity pixel [11]

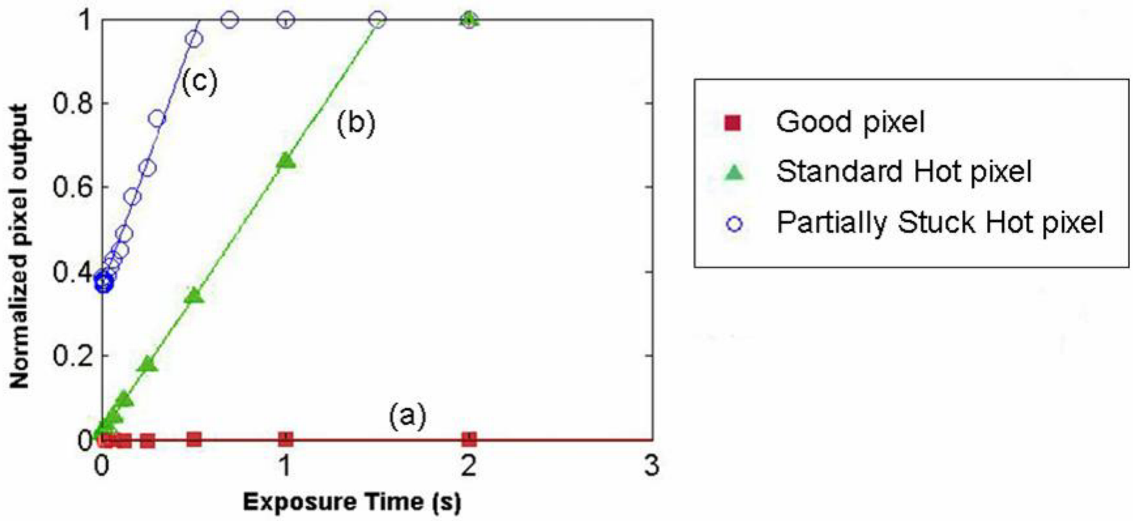

Fig. 9: Hot pixel response [15]

defect. According to Leung et al. [17], the majority of hot pixel defects found in larger sensor area DSLRs are low impact defects with a low dark current level. Therefore, these defects are almost invisible in images captured under normal conditions (short exposure time and low ISO level) and have only little influence on the image quality.

A typical method to compensate hot pixels used by many camera manufacturers is to capture a dark frame of the same exposure time right after the image is captured and subtract this one from the original image, which should eliminate the dark signal then. An interesting fact is that hot pixels become more visible if many images are captured in rapid succession while the response of good pixels remains the same. This might be due to heating of the image sensor, which leads to an increased dark current and therefore a higher hot pixel amplitude. Thus the automatic detection of hot pixels gets even more challenging because of the varying dark signal and therefore varying hot pixel amplitude.

3.4.10 Partially-Stuck Hot Pixel

The literature is not completely clear about how to name this defect. According to Leung et al. [14], a partially stuck hot pixel has an additional component (offset) that is independent from the illumination and exposure time. Therefore, it can even be observed at no exposure (0 exposure time) and will in general appear brighter than a standard hot pixel. Thus it has a higher impact on the image quality because the pixel goes into saturation at a lower illumination level, which further reduces the dynamic range of the pixel. In addition partially-stuck hot pixels cannot be compensated by the simple method described above for standard hot pixels, because if the combined illumination and dark signal causes the pixel to saturate, just by subtracting the dark response will not recover the original signal. Jessica Fridrich [18] does not explicitly mention partially stuck hot pixels.

3.5 Defects in Colour Images

Unless images are captured in RAW mode, several imaging functions are applied to the output images captured in colour mode, which is the usual mode used in photography. These imaging functions include demosaicing, noise reduction, white balance etc. Although they are intended to improve the image quality, they ignore the presence of defects, i.e. defective pixels are treated as good ones, so they might not only amplify but also spread the impact of a defective pixel to its neighbouring good pixels. Especially demosaicing, which is the first function applied in the imaging pipeline, has a significant impact on the appearance of defective pixels. It is needed for every CFA image sensor to recover the missing two colour channels at each pixel. This is done by interpolation using the neighbouring pixels but ignoring the presence of defects, which will lead to incorrect interpolation values. Depending on the type of interpolation used, it will distort the defective pixel, causing a singe defect to appear as a virtual cluster. Such a defect cluster is more visible than a single pixel defect. Thus sensor defects are highly undesirable in colour images. Nowadays the most widely used demosaicing algorithms are adaptive ones, taking the characteristics of the image into account, similar to the Kimmel algorithm, which reduces the Moire pattern. But these more advanced demosaicing algorithms also have a higher impact on the defect appearance compared to simple bilinear interpolation and may even spread defects into other colour channels. To be able to detect the defect location, each local cluster is considered as one defective pixel and the peak of the defect cluster appearing in the output image is used as defect location.

Most cameras use JPEG compression for colour images. This compression reduces the peak error, but may spread a single defective pixel or a virtual defect cluster even wider, which of course has a negative effect on the output image. JPEG compression also suppresses colour variations of all the three colour planes which lowers the impact of the defective pixels.

All these imaging functions applied to colour images have an irreversible impact on the pixel defects in colour images, in the way that they further degrade the quality of output images and make the determination of the exact defect location more difficult, but the defect itself may be easier to detect because its higher impact. Therefore, a suitable defect detection algorithm has to be adopted to be able to handle colour images correctly.

3.6 Defect Identification Techniques

There are different techniques which can be used to identify the different types of defective pixels. With DSLR cameras all of these techniques can be used, because these cameras do not only provide the RAW format, but also manual control of camera parameters, especially exposure control. If PS (point-and-shoot) or mobile phone cameras should be investigated some of these techniques cannot be used at all or have to be adopted.

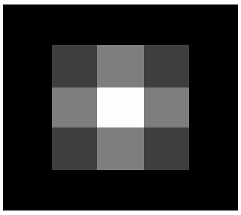

Fig. 10: Defect cluster after demosaicing [14]

Fig. 11: Red hot pixel in RAW image (left) and after processing (right) (from: `http://www.pixelfixer.org/faq.html`)

3.6.1 Dark Field Calibration

The dark field calibration can be used to identify stuck-high, fully-stuck, hot and partially-stuck hot pixels. Standard dark field calibration uses a dark field exposure, i.e. taking an image with no incident illumination. Therefore the output image should be completely black and a stuck high or hot pixel would appear as bright dot. Dark field calibration can be easily done even with digital cameras not having a shutter. A stuck high or a hot pixel's location can be identified using a simple threshold test to distinguish a sensor defect from random noise, e.g. if the output of the suspect pixel is 3 times higher than the noise signal it is declared as defective. Figure 12 shows an example of a dark field calibration image.

The distinction between stuck high and hot pixels can be made by capturing several dark field images with increasing exposure times. According to the pixel model of Leung et al. [13, 14], a hot pixel has an illumination-independent component that increases linearly with exposure time, so it will appear brighter if the exposure time is increased, while a stuck pixel does not change its intensity value.

The dark response I_{offset} of a hot pixel, also called combined dark offset, can be estimated by setting $R_{photo} = 0$ in the pixel model (section 3.4.2), then it becomes

$$I_{offset}(R_{Dark}, T_{exp}, b) = m \cdot (R_{Dark} \cdot T_{exp} + b)$$

which is independent of the incident illumination. A linear function can be used to estimate the defect parameters R_{Dark} (slope of the pixel's response) and b (offset along the y-axis of the pixel's response), see figure 9 for details. For a standard hot pixel $b = 0$ while for a partially-stuck hot pixel an additional, illumination and exposure independent offset is added and $b \neq 0$. Standard hot pixels

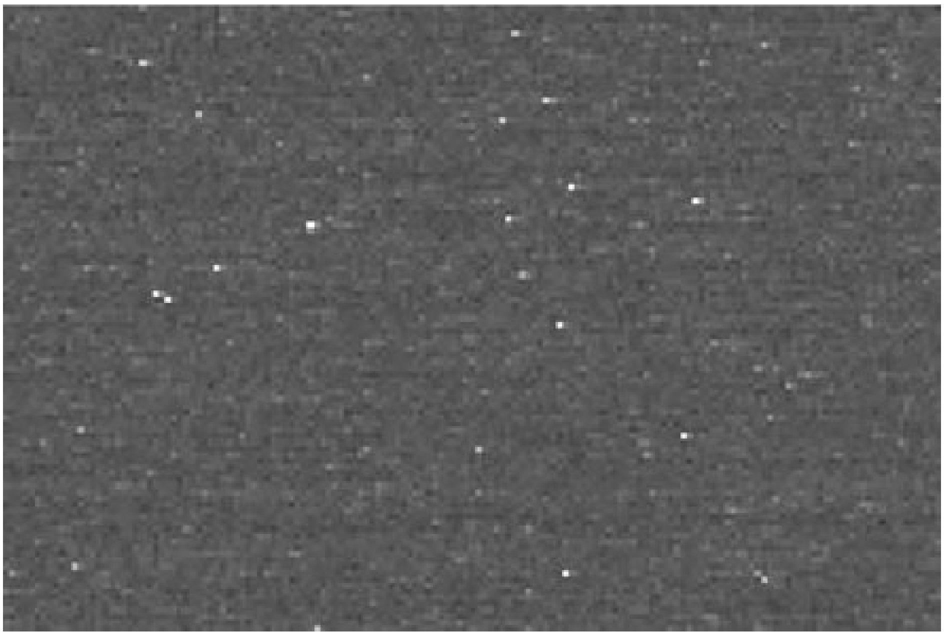

Fig. 12: Dark field capture example [18]

are most visible in images taken with long exposure times, while partially-stuck ones are visible in all images depending on the offset value b. The hot pixels are then identified using a threshold test to eliminate sensor noise.

3.6.2 Bright Field Calibration

Similar to the dark field calibration, in bright field calibration a bright field exposure is used. This technique can be used to identify stuck-low defects and also fully-stuck defects. Dudas et al. [12] capture an image with uniform illumination near saturation. In the output image all good pixels should have the same intensity value, whereas stuck low pixels would appear darker. Like in a dark field image, also here the pixel defects can be detected using a simple threshold test. According to Leung et al. [16] they never found a single stuck low pixel (except for manufacturing time defects which are normally masked out). Jessica Fridrich [18] suggests to capture an image of a partly cloudy sky instead. In such an image also stuck high pixels are visible if they are stuck on a very high value near saturation. Especially for bright field calibration a set of calibrated equipment is necessary to reproduce a uniform illumination at a known intensity, which is necessary to estimate the defect parameters in addition to the defect location. So the complexity and high costs make these defect detection method impractical for in-field use and thus it is typically used by image sensor manufacturers only.

3.6.3 Defect Identification for PS and Cellphone Cameras

Leung et al. [14] pointed out that the dark field calibration procedure as described above cannot be used for PS and mobile phone cameras due to the limited exposure control. In fact the camera exposure system has to be tricked to give maximum exposure time, which is necessary to be able to do a dark frame calibration at all, but there is no ability to use increasing exposure times. Thus it cannot be distinguished between a stuck-high and a hot pixel. Moreover these cameras do not provide the

RAW format, therefore the calibration images are captured in colour mode. As described in section 3.5, in colour mode several imaging functions are applied (main problem is the colour demosaicing), therefore a single defect will appear as a virtual cluster. Thus each local cluster has to be considered as one defective pixel using its peak output as the defect location. Each time they captured a series of 3 images and used the common result to eliminate false detection due to noise.

Another problem regarding mobile phone cameras is that for cost reasons there is no factory calibration done. Therefore it cannot be distinguished between factory time and in-field defects at the first measurement (initial number of defects may include factory time defects) but only on subsequent measurements. This is also the reason why there might not only be hot pixel defects but also stuck and low sensitivity pixel defects found across these types of cameras.

3.7 Impact of Sensor Parameters on Defect Growth Rate

Chapman et al. [9, 6, 7] showed by continued empirical measurements of several imagers with different characteristics (DSLRs, PS cameras and mobile phone cameras) that the rate of newly generated pixel defects depends on the sensor technology (CCD or APS) and on sensor design parameters like sensor area, pixel size and gain (which is adjusted using the ISO value), pointing out important trends in defect development. They showed that the defect rates scale linearly with the sensor area, so they introduced a metric called defect density measured in defects/year/mm². The dependence of the defect density on pixel size and ISO is not linear. Instead they found the best fit by using a power law relationship, according to:

$$D = A \cdot S^B \cdot ISO^C \tag{2}$$

where D is the defect density (defects/year/mm²), A is the number of defects/year/mm² if the pixel size is $1\,\mu m$, S is the pixel size in μm, ISO is the ISO level and B and C are constants depending on the sensor type. For a CCD sensor $A = 0.0141$, $B = -2.25$ and $C = 0.69$ and for an APS sensor $A = 0.0742$, $B = -3.07$ and $C = 0.5$.

3.7.1 Sensor Type

In general, CCD sensors tend to have a higher defect count than APS ones of the same age ($\sim$ 3x higher than APS for DSLR sensors with the same size at all ISO levels). If cosmic ray radiation is the main source of the defects, then an explanation why CCD sensors are more sensitive to the development of in-field defects is that on the one hand the size of a CCD sensor pixel is 2x larger than the size of an APS pixel and on the other hand CCD sensors have a higher fill factor (70-90%) than APS sensors ($\sim 25 - 30\%$). This means that at the same pixel size, the photosensitive area of each pixel is only about 25% for an APS sensor compared to about 80% for an CCD sensor. As the photosensitive area increases and because the cosmic ray total flux is constant per area, the defect rate is also increased.

But this is only true for pixel sizes of about $6 - 7\,\mu m$. For a pixel size of $2.1\,\mu m$, APS and CCD sensors would have the same defect density. For even smaller pixels, APS sensors would have a higher defect rate than CCD ones. This means that CCD pixels are not as sensitive to a shrinkage in pixel size as APS ones, which can also be seen from the value of B in the formula. The impact of the pixel size on the defect density is described below.

3.7.2 Sensor Area

The linear increase in the defect density with increasing sensor area can also be explained based on cosmic ray radiation as the main source causing the defects. According to [3, 4] the cosmic ray total flux per area is a random process with only minimal changes over time in a given location, i.e. it scales with the area and therefore a larger photosensitive area increases the exposure to the radiation. But the defect rate only scales linearly with the area if the pixel size remains the same. Otherwise the change in the defect density also depends on the change in pixel size as described below.

3.7.3 Pixel Size

The pixel size has a high impact on the defect density. In the last few years the sensor sizes have remained the same, but the number of pixels increased. This can only be achieved by reducing the pixel size. Especially sensors equipped in mobile phones have very small pixel sizes (side length of 1-2 µm). If the pixel size would have no impact, one would expect the sensor with the highest area to develop the most defects if the pixel size is kept constant. But instead the expected rate from scaling up the defect rates of PS and cellphone cameras by the sensor area is much higher than the observed rates for the DSLRs. As the equation 2 shows, the defect density increases rapidly with decreasing pixel size. It scales approximately inverse linear with the pixel area for CCD sensors and much more than inverse linear with the pixel area for APS sensors, which means that sensors with smaller pixels are more sensitive to developing defects. As the relationship between pixel size and defect rate is exponential, the impact of reducing the pixel size is much higher than increasing the sensor size. As most of the smaller sensors are APS ones, this has a great impact on the quality of the output images. If the pixel size shrinks down to 1 µm the increasing defect rate will require a significant amount of defect suppression and even the usual mapping done by the manufacturers to ignore the output of hot pixels and replace their values by interpolation from other pixels may fail as a defect correction strategy.

According to Chapman et al. [5] one reason for this much higher defect rate might be that due to the smaller area per pixel only a smaller fraction of the incoming light (about 10% compared to DSLRs) is collected. Therefore to create a similar output, the scaling factor m in the pixel model has to be increased. This is similar to increasing the ISO level as described below, therefore also the impact of the defects, especially I_{Offset} increases. This means that hot pixels which would have a low influence on a sensor having a larger pixel size are much more visible at this small pixel size. Another reason is that the dark current of a hot pixel, which is caused by the defect source, remains the same as the pixel size shrinks, but the sensitivity of the pixel to each electron increases and therefore even a small hot pixel damage will cause a significant effect on a smaller pixel. Currently the reason for the difference in sensitivity between CCD and APS sensors if it comes to pixel size is not completely clear and needs to be explored.

3.7.4 ISO Level

In previous times the usable ISO range was limited by the sensor noise. But as sensors got better, the noise level decreased and in addition noise suppression algorithms are utilized, pushing the usable ISO ranges for modern DSLR cameras up to 25600 now. Capturing images at high ISO levels enables natural light photography without the need for a flash or resorting to long exposure times which is a great benefit for photographers. But unfortunately the pixel defects are increasing in their number with higher ISO values. Leung et al. [17] showed this by performing the dark frame calibration at various ISO levels. The reason is that the ISO setting is nothing else than a numerical gain (see the pixel model in equation 1), thus also the defect parameters (offset and dark current) are amplified (scaling linearly with the ISO level), enhancing the brightness of a hot pixel. As the offset is amplified, partially-stuck hot pixels are going into saturation even at short exposure times, making them appear like a stuck high defect. This is also a major drawback for the pixel's operation, because this high offset reduces the dynamic range of the pixel. Of course the threshold for determining a defect also has to be increased as the sensor noise increases. But the brightness of the defects increases at a much higher rate than the background noise signal, therefore at higher ISO levels more defects can be

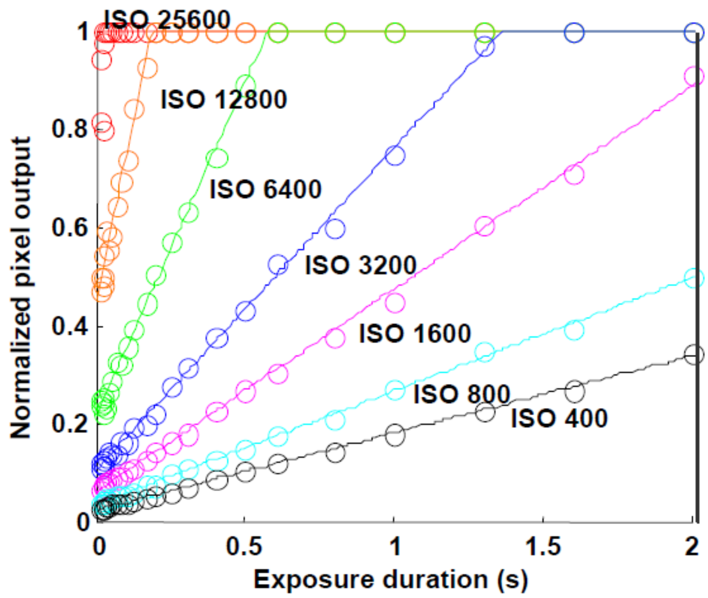

Fig. 13: Dark response at various ISO levels [9]

found, because they are more distinguishable from the noise. Leung et al. [19] noticed a substantial increase in the number of defects detected at ISO 3200 compared to ISO 100. The reason for this increase is that many of the defects are barely visible at lower ISO settings due to low R_{Dark} and b values, i.e. they cannot be distinguished from the noise signal, so they may not pass the threshold test. The defects present at higher ISO levels affect a wider range of exposure times and are visible even at short exposure times (due to their high offset). According to the quantitative values, CCD sensors are more sensitive to an increase in the ISO level than the APS ones.

4 Defect Detection Algorithm

Each image is a snapshot of the sensor's state at a given time, i.e. of all of its pixels, including the defective ones. As described above, defects are usually detected using special calibration images (dark and bright field exposure). Unfortunately such images were not available for the type of fingerprint, finger vein and hand vein sensors utilised to capture the databases I worked with during my evaluations. Using calibration images the incident illumination is uniform and known, therefore the defects can be easily detected and their parameters can be estimated. If only fingerprint, finger- or hand vein images, respectively, captured by a specific sensor are available, the illumination is neither uniform nor known. Moreover, factors like the ISO level, exposure time and the complexity of the image scene will affect the visibility of the defects. In this case, some kind of statistical approach using a sequence of images has to be used to identify the defects.

Of course all of the methods described below are only heuristics and will never achieve an accuracy of 100%, i.e. identify each defect correctly. Thus the number of defects and their parameters detected by these methods can only be an approximation of the real values. The majority of photographs are taken at short exposure times ($\sim 1/30\,s$) and low ISO settings (<ISO800) [17], thus low impact defects, especially hot pixels with a low dark current magnitude, are barely visible in these images and might not be detected at all. Long exposure times are not used often, because camera motion will cause blur in the image unless a tripod is used.

The EXIF data provides additional information (exposure time, ISO setting, capture date), which can be used to estimate the parameters of the defects (dark current amplitude, offset, etc.) but unfortunately this information is not available for most fingerprint, finger vein and hand vein images, because these are usually stored in lossless formats like Bitmap (bmp), Portable Network Graphics (png) or TIFF which do not contain EXIF data.

4.1 Filters for Defect Identification

Sensor ageing related defects occur in the form of single pixel, point like defects as discussed above. Therefore filters like median, ring averaging and 4NN and 8NN minimum distance are suitable to detect the defects. The image is filtered, then the actual pixel value is compared with the one in the filtered image. The probability for the centre pixel to be defective corresponds to the difference of the pixel values.

4.1.1 Median Filter

The median filter is defined by:

$$y(i,j) = \underset{(m,n)\in G(i,j)}{median}\{I(m,n)\}$$

where the pixel (i,j) is the centre pixel which is under examination, $y(i,j)$ is the output of the filter, $G(i,j)$ denotes the locations of pixels within a neighbourhood around the centre pixel. This can be any type of neighbourhood. $I(m,n)$ are the outputs of the neighbouring pixels.

4.1.2 $n \times n$ Ring Averaging Filter

In contrast to the standard averaging filter, which takes the average of all the pixels inside a square or rectangular neighbourhood around the centre pixel, this filter only takes the average of all pixels inside a ring, centred at the centre pixel as output value. This mitigates the effect of defect spreading due to colour demosaicing, which mostly affects pixels directly neighbouring the defective centre pixel, as only pixels on an outer ring are considered for the calculation of the interpolation value.

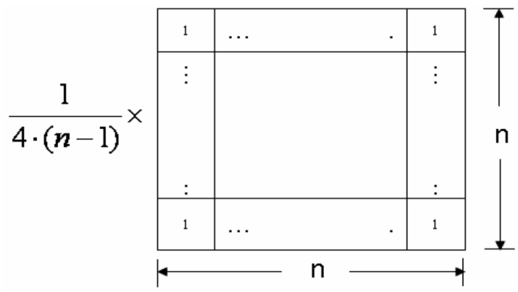

Fig. 14: $n \times n$ ring averaging filter [16]

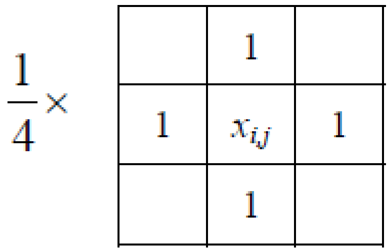

Fig. 15: 4NN filter [16]

4.1.3 4NN Filter

This filter basically takes the mean of the centre pixel's north, south, east and west neighbour as output value.

4.1.4 4NN and 8NN Minimum Distance Filter

The minimum neighbouring pixel difference filter can be defined by:

$$y(i,j) = \min_{(m,n)\in G(i,j)} \{|I(i,j) - I(m,n)|\}$$

where the pixel (i,j) is the centre pixel which is under examination, $G(i,j)$ denotes the locations of pixels within a neighbourhood around the centre pixel. This can be an 8 neighbourhood (3x3 window centred at the pixel and using all pixels without the centre one) or a 4 neighbourhood (using only the pixel directly south, north, west and east of the centre one) or any other type of neighbourhood. $I(i,j)$ is the output of the centre pixel and $I(m,n)$ are the outputs of the neighbouring pixels.

4.2 Thresholding Based Approach

This is a two step approach. At first the stuck pixels are identified using the same property as in the statistical approach proposed by Bergmüller et. al [22], which is described below, but no hard comparison (a stuck pixel is supposed to have exactly the same value in all images) is used. Instead a certain threshold can be set (in percent of the maximum pixel value) to account for sensor, analog to digital conversion and transmission noise, respectively, which means a stuck pixel might not always have exactly the same pixel value. The second threshold defines in how many images (also in percent of the total number of images) the pixel must have the same value within the tolerance to be identified as stuck pixel.

In [24] a pixel is identified as low sensitivity pixel if:

$$|\frac{\sum_{k=1}^{N} P(k)}{N} - P(0)| \geq C, \; C = 10\%$$

where $P(k)$ describes the value (output) of a pixel, the pixels $k = 1...N$ describing a neighbourhood (e.g. window centred at the pixel) around the pixel and $P(0)$ is the output of the centre pixel. $C = 10\%$ means 10% of the maximum possible output value of a pixel. A pixel is identified as stuck low, i.e. it has a value of 0 if:

$$|\frac{\sum_{k=1}^{N} P(k)}{N}| \geq C_{011}$$

$$P(0) \leq C_{012}$$

where C_{011} and C_{012} are 50% and 10% of the maximum possible output value of a pixel, respectively. This means that a pixel is only identified as stuck low if it has a value near zero and its neighbouring pixels have a value greater than or equal to127 (for 8 bit images) on average. A pixel is identified as stuck high if:

$$|\frac{\sum_{k=1}^{N} P(k)}{N}| \leq C_{101}$$

$$P(0) \geq C_{102}$$

where C_{101} and C_{102} are 90% and 50% of the maximum pixel value, respectively. This means that similar to stuck low pixels, a pixel is only identified as stuck high if its value is near the maximum value and its neighbouring pixel values are below or equal to 127 on average. Of course the threshold values C can be adjusted according to the input images. In the original paper only a single image is used, not an image sequence. This approach was also tested for the implementation with an image sequence but did not improve the result, so the simple method using the two thresholds was used.

The next step is the identification of the hot pixels. A hot pixel is a point defect, i.e. a single pixel, which appears brighter than its neighbouring pixels, so it can be identified by comparing the pixel's value to the value of its neighbouring pixels. Therefore, the filters described above can be used to generate an interpolation (prediction) image and compare the prediction with the actual pixel value then. If the error is above a certain threshold, the pixel is identified as hot pixel in the current image and again if this pixel is identified as hot one in a certain percentage of all images (above the second threshold) it is finally identified as hot pixel. This approach is only able to estimate the number of hot pixels but not their parameters.

4.3 Statistical Approach

This approach was proposed by Bergmüller et. al [22] to estimate the defect growth rates and amplitudes using two image sets, taken with the same sensor at different points in time. $Y_0...Y_K$ is a sequence of K images taken in a very short period of time. A stuck pixel has the same value in each image of this sequence, therefore a pixel at a given position is stuck if:

$$y_0 = y_1 = ... = y_K$$

A partially stuck pixel[1] at position (x, y) adds a light independent offset to the pixel's output at that position $y_k = Y_k(x, y)$. First the pixel's mean $\bar{y}$ is computed from the whole image sequence. The mean is then substituted using the pixel model from section 3.4.1. Because stuck pixels can be detected using the above method, the possibility of a pixel being stuck can be ruled out, thus c is not considered here. To take impredictable impacts during acquisition into account, some modelling noise is added, leading to the following equations:

$$\bar{y} = \frac{1}{K}\sum_{k=1}^{K} y_k$$

$$\bar{y} = \frac{1}{K}\sum_{k=1}^{K}(i_k + d + \theta_k)$$

If K is sufficiently large, the mean of a uniformly distributed modelling noise θ cancels out. Under the assumption that there are no pixel defects present, the offset $d = 0$ which is denoted as d'. Under this assumption $i_k = y_k$ can be substituted and the pixel's mean $\bar{y}$ is denoted as $\bar{y}'$. The partially-stuck (I denote it as hot pixels) defect matrix d contributes to each pixel's output uniformly, thus it is independent of k. If there are no pixel defects and if the pixel means $\bar{y}$ mostly have identical values in a $q \times q$ neighbourhood, the identity $\bar{y}' \equiv median(\bar{y}, q)$ holds. This is the case if only pixel means of image regions with mostly uniform brightness and texture, e.g.. outside the finger area in fingerprint images, are considered. Thus regions covered by the finger area have to be masked out (Bergmüller et al. [22] originally masked out the area where the iris texture is). As the defects are distributed uniformly across the sensor area (i.e. inside the image) it is sufficient to use only a certain part of the image, search for defects in this area and derive the total number of defects inside the whole image based on the percentage of the whole image area which is covered by the part used to search for defects.

Sparse outliers within the $q \times q$ region are filtered out due to the non-linear nature of the median filter. Such outliers occur if there is a pixel defect. Thus the median is able to compute a pixel's mean $\bar{y}'$ disregarding the influence of pixel defects. Using this relation, $\bar{y}' \to median(y, q)$ is set and thus $d' \to \hat{d}$ can be used as an estimator for the offset, which is added to the incident illumination i_k, leading to the pixel's output y_k.

$$\bar{y}' = d' + \frac{1}{K}\sum_{k=1}^{K}(y_k)$$

$$\hat{d} = \frac{1}{K}\sum_{k=1}^{K} y_k - median(\bar{y}, q)$$

Not only information about the partially-stuck pixels but also information about the PRNU is contained in $\hat{d}$. The PRNU is likely to be normally distributed, therefore the size of the median kernel q has to be chosen large enough to minimize this influence on the $median(\bar{y}, q)$, which is the case if a normal distribution in the logarithmic histogram is observed inside the $q \times q$ neighbourhood, like it is the case for the whole image.

A pixel is identified as partially-stuck if $\hat{d}$ is an outlier with respect to the normal distribution, i.e. it has a much higher sensitivity than expected from the PRNU. This is done using a decision threshold $\hat{d} > \tau_{ps}$ which is chosen manually.

[1] It is usually denoted as hot pixel during this thesis. But to comply with their original explanation it is called partially stuck pixel in the following.

4.4 Bayesian Inference Based Approach (Dudas)

This is another statistical approach but in contrast to the one above it uses Bayesian inferences to calculate the probability of each state for each pixel separately from a series of images taken by the sensor within a short timespan. It is based on the defect model described in section 3.4.2, which models an imager as an array of $W \times H$ pixels where the illumination present at a pixel at position (i, j) is denoted as $x_{i,j}$. The output (response) of each pixel is denoted as $y_{i,j}$, is in the range of $[0, 1]$ and can be written as a linear function of the incident illumination:

$$y_{i,j} = m_{i,j} \cdot x_{i,j} + b_{i,j}$$

where the subscripts are omitted for simplicity. m is the sensitivity or gain and b is the bias or offset, both having characteristic values for specific defect types (see pixel model in section 3.4.2). An ideal, non-defective pixel would have $m = 1$ and $b = 0$. Each defect type can therefore be described by a combination of gain m and offset b and is referred as $D = (m, b)$. A stuck high pixel can be described by $(0, 1)$ and a stuck low one by $(0, 0)$. The probability of the occurrence of a specific defect type (m, b) is denoted as $p_{(m,b)}$ and the total probability of any defect occurring, i.e. the defect density, can be calculated by $p_{total} = \sum_{\{D\}} p_{(m,b)}$. To completely calibrate an image sensor, the parameters m and b, i.e. the sensitivity and bias of each pixel has to be determined.

As the defects cannot be reliably determined solely based on one image, a sequence of images has to be used. These are images captured in the field under normal operation conditions which are analysed and the statistics of each image and defect type are accumulated using Bayesian inferences. The details of the algorithm are described in [10, 11]. After capturing $T + 1$ images the samples $Y = y_{i,j}^{(0)}, y_{i,j}^{(1)}, ..., y_{i,j}^{(T)}$ from a specific pixel at position (i, j) are available. Using Bayesian inferences the probability of each defect type D at each pixel position (i, j) can be estimated based on the collected evidence (image samples) using (subscripts are omitted for simplicity):

$$P(D|Y = y^{(0)}, y^{(1)}, ..., y^{(T)}) = \frac{P(Y|D) \cdot P(D)}{\sum\limits_{D'=\{all\ D\}} P(Y|D') \cdot P(D')}$$

where the conditional probability $P(Y|D)$ is the most important term, which encapsulates the probability of observing the data Y under the assumption that the given defect type D is present at that pixel. This likelihood has to be evaluated within the context of the image being processed. The Bayesian function accumulates the statistics over a sequence of images to estimate the probability of a pixel to be of a certain defect type D. There are two schemes for the evaluation of this data based on statistical metrics calculated from the image, which are described in the following subsections. $P(D)$ denotes prior knowledge of the occurrence probability of each defect type and can be obtained from manufacturing test data. These calculations are done recursively, such that each image in the sequence is processed individually. The previous iterations lead to the a priori probability at iteration k. This is expressed by the following formula for the k-th iteration:

$$P(D|Y = y^{(0)}, y^{(1)}, ..., y^{(k)}) = \frac{P(y^{(k)}|D) \cdot P(D|y^{(0)}, y^{(1)}, ..., y^{(k-1)})}{\sum\limits_{D'=\{all\ D\}} P(y^{(k)}|D') \cdot P(D'|y^{(0)}, y^{(1)}, ..., y^{(k-1)})}$$

This conditional likelihood term is evaluated for the whole sequence of T images and after that, the defect type with the highest likelihood at a given pixel is taken as resulting defect type. Of course the set of all defect types should also include good pixels to be complete. The parameter T is set before the testing begins. T should not simply be chosen to use all available images, because long sequences may cause the statistics to saturate. Other termination criteria, like a pre-determined confidence level, i.e. likelihood threshold are also possible. Then the convergence speed is strongly affected by the content of the images in the sequence. A flowchart of this approach is shown in figure 17.

Especially if not only the defect location but also the earliest date in time where the defect showed up should be determined (which can be done if the image sequence covers a longer timespan), a sliding window approach should be used, at which only the n most recent images are used to calculate the accumulation. According to Leung et al. [15], a window length of 3 is best for accurate temporal detection of low visibility defects. For an accurate detection of the date when the defect occurred first in time, a dark-frame calibration image is essential, from which the first prior probabilities $P(D)$ for each pixel and defect type can be derived. Changes in the accumulated probability indicate the presence or absence of defects. Therefore, the date when a pixel first got defective can be either determined using a simple threshold test, i.e. if $P(Good|y^{(x)})$ falls below a certain threshold or using a probability ratio threshold test: $P(Good|y^{(x)})/P(Good|y^{(k-1)}) > threshold$.

A problem regarding this approach might occur if only a certain type of images, e.g. landscape scenes with large black and underexposed areas are used. Then defects like stuck low pixels cannot be detected. In fingerprint images where the outer image regions are rather bright all the time, defects like hot pixels might remain undetected if their amplitude is rather low. Finger-vein images are usually dark grey to black outside the finger region, consequently the detection of stuck low pixels may not be possible. Also hand vein images may exhibit such low brightness regions if not only a ROI but the whole image of the hand including the surrounding areas is used. According to Dudas et al. [12] this method can only be used to detect hot pixels if images with sufficiently long exposure times are available.

4.4.1 Image Statistics Method

This technique relies on image wide statistics to derive the probability for each defect type and is therefore called image statistics method [25]. It relies on the image-wide distribution of pixel values in the form of the density function $P_X(x)$, which is estimated for each image using the image histogram $P_Y(y)$. E.g. the likelihood $P(y^{(k)}|D)$ for a pixel to be stuck at a value c simply is:

$$P(y^{(k)}|D(0,c)) = \begin{cases} 1 & y^{(k)} = c \\ 0 & otherwise \end{cases}$$

and for a pixel with half sensitivity it is:

$$P(y^{(k)}|D(0.5,0)) = P_X(2y^{(k)})$$

The problem of this simple approach is that a stuck pixel might not always have the same output due to noise. Thus the conditional probability models have to be adjusted in a way that the statistical distributions do not represent a single impulse function but a smooth function, like a Gaussian distribution with its mean value at the stuck pixel value and a small variance. The likelihood for a pixel being good is given by $P(y^{(k)}|D(1,0)) = P_X(y^{(k)})$. If the output for a specific defect type gives unlikely values after testing many images, that defect type is not very likely to be present at that pixel. Therefore, pixels that are consistently having an output of 0 or 1 are not very likely to be good, because these extreme values are not likely to occur very often in typical images. But if only global information is used and local information around the pixel of interest is not taken into account, this may lead to false identifications depending on the images contained in the sequence. E.g. in the context of fingerprint images the outer region might always be rather bright depending on the type of sensor, which might lead to the false identification of stuck high pixels. Finger- and hand vein images show rather large bright parts and only little dark parts where the veins actually are. Outside the hand or finger region there are some quite dark parts. This may lead to both, the false detection of stuck-low and stuck-high pixels inside the respective parts. If the information of neighbouring pixels is additionally used, these false identifications can be avoided, because it is unlikely that not only a single pixel is stuck high but also all of its neighbours are. In a photograph saturated white areas are undesirable, therefore the dynamic range of the image is limited by the image processing chain of many cameras, which could affect the image statistics $P_X(x)$ and $P_Y(y)$ in undesirable ways. This leads us to the interpolation technique which is described next.

4.4.2 Interpolation Method

The more successful technique for calculating $P(Y|D)$ or $P(y^{(k)}|D)$, respectively uses information from the pixel's neighbourhood. The details of this approach are described in [11]. The output of each pixel is compared to its neighbours and the probability that this pixel is faulty is calculated, assuming that most of the pixels in the neighbourhood of a defective pixel are good ones. This is done by first estimating the incident illumination at the pixel's location (i, j) by interpolating its neighbouring pixels' values. Of course the choice of the interpolation scheme has a significant impact on the accuracy of the detection method. Although several interpolation schemes can be used, for this application the emphasis is placed on different performance parameters than for typical applications. E.g. many schemes try to maintain high-frequency components in the output and try to use as much of the surrounding data as possible, which is both contra productive for defect detection. The authors proposed the 4NN and ring averaging schemes, which are shown in the figure 16. The smaller the interpolation region and the closer this region is located towards the pixel under investigation, the better will be the estimate of the pixel's value. According to Jenny Leung [21], the 5x5 ring averaging filter provides the best tradeoff between accurate pixel estimation and avoiding the defect spreading problem caused by colour interpolation if dealing with colour images. If the estimated pixel value at position (i, j) is denoted as $z_{i,j}$, then the estimation error can be calculated by $e_{i,j} = z_{i,j} - y_{i,j}$ for each pixel in the image. Based on the estimation error, the estimation error density PDF $p_E(e)$ is calculated over the whole image and its statistics are accumulated in the probability distribution (CDF) $P_E(e)$, which also describes the effectiveness of the estimation scheme. Now the likelihood for each defect type can be determined by calculating the expected error e_D at each pixel, for each defect type D and retrieving the probability for this defect type from the distribution $P_E(e_D)$ using

$$P(y^{(k)}|D) = P_E(e_D^{(k)} = f_D(z) - y^{(k)})$$

$$P(y^{(k)}|(m, b)) = P_E(e_D^{(k)} = z - [\frac{y^{(k)}}{m} - b])$$

where f_D describes the transfer function of the defect type. E.g. a pixel stuck at value c has $P(y^{(k)}|D(0, c)) = \begin{cases} 1 & y^{(k)} = c \\ 0 & otherwise \end{cases}$ and a a half sensitivity pixel has $P(y^{(k)}|D(0.5, 0)) = P_E(z - 2 \cdot y)$. Like before, the predicted defect type at the pixel (i, j) is the one with the highest likelihood then. The estimate from the interpolation does not need to precisely predict $x_{i,j}$ at each image, because using a certain number of test images, the results should be reasonably accurate. Thus this scheme is more effective than the image statistics method. The authors claim that only 50 ordinary images are sufficient to accurately identify all faults without falsely diagnosing good pixels as faulty [11]. They also found out that the 8NN or 3x3 averaging filter work better than the 4NN averaging filter because due to the smaller weighting of each single pixel defects have less impact on the interpolation results. Furthermore adding noise only slows down convergence rates but does not affect the false positive detections. Using other interpolation schemes like the 4 round-robin technique, which is a modification of the 4NN scheme and described in [10] further mitigates the effects of a catastrophic

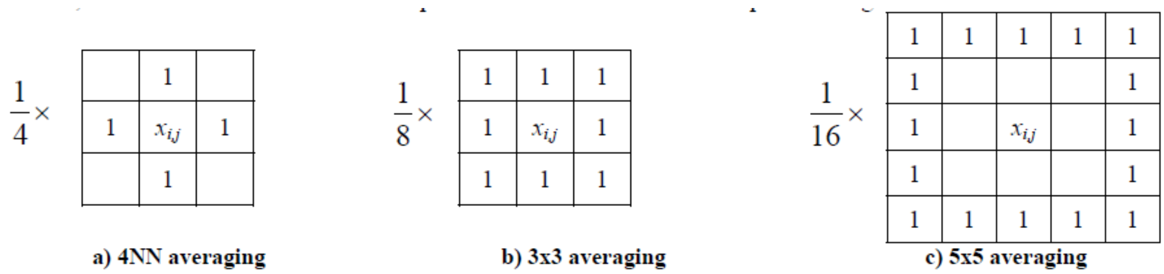

Fig. 16: Filters for interpolation method [15]

error but might be less accurate. It basically splits the 4NN interpolation into 4 separate tests and uses a majority voting scheme. At first the pixel is compared to its left neighbour and the probability $P(D|y^{(k)})$ is evaluated, then it is compared to the neighbour above, again evaluating the probability using the left pixel's test probability as prior probability. This is continued until all 4 neighbours have been examined. If the interpolation error becomes very large this causes issues within the Bayesian statistical methods (saturation), which leads to false positive detections and should therefore be avoided.

As mentioned before if only colour images are available for defect detection, the single pixel defects are distorted by the multiple steps in the image processing chain of the camera to produce a colour image from the raw data, including demosaicing, white balance, noise reduction, linearisation, image scaling, dark signal removal, exposure compensation and compression. All these imaging functions will make a single pixel defect look like a defect cluster, which cannot be detected by the simple interpolation schemes any longer. A possibility to deal with this is to use a 5x5 or 7x7 ring averaging filter, because the pixel's nearest neighbours are affected most by the colour demosaicing. Thus interpolating using pixels further away can mitigate this effect, but also leads to a less accurate interpolation, i.e. a higher interpolation error, because of the larger interpolation region. As a consequence the detection result might not be very robust. The defect location might be detected (actually easier because it is more visible), but the defect parameters cannot be detected reliably. Thus access to the raw data is favourable for an accurate defect detection.

Another issue might be if the defect density gets rather high (5% of all pixels are defective). Then the performance of this scheme suffers because the distribution of $P_E(e)$ is derived from an image with defects and the more defects the less accurate this CDF becomes. Another reason might be that some defect locations are better to detect defects than others, depending on the surrounding pixels' values. The more defects there are the higher is the probability of defects to be in these regions and thus it gets more difficult to accurately detect them. The performance of this scheme reduces dramatically if multiple defects within a 3x3 pixel area are allowed, because of the interpolation technique used. The accuracy of the estimate of $P(D|Y)$ highly depends on the estimate of $z_{i,j}$, which is calculated by an interpolation of the neighbouring pixels' outputs. If these pixels are also defective, the interpolated values will be incorrect, which leads to an incorrect estimate of the probability $P(D|Y)$. These problems illustrate the shortcomings of using local information from a pixel's neighbourhood to detect defective pixels. They also illustrate the importance of statistically evaluating data instead of using fixed thresholds and explain why statistical methods are more reliable. A large interpolation error in combination with fixed thresholds may lead to a large number of false positive detections using fixed thresholds, but due to the Bayesian inferences these effects might be mitigated or at least be not as catastrophic when using this statistical approach.

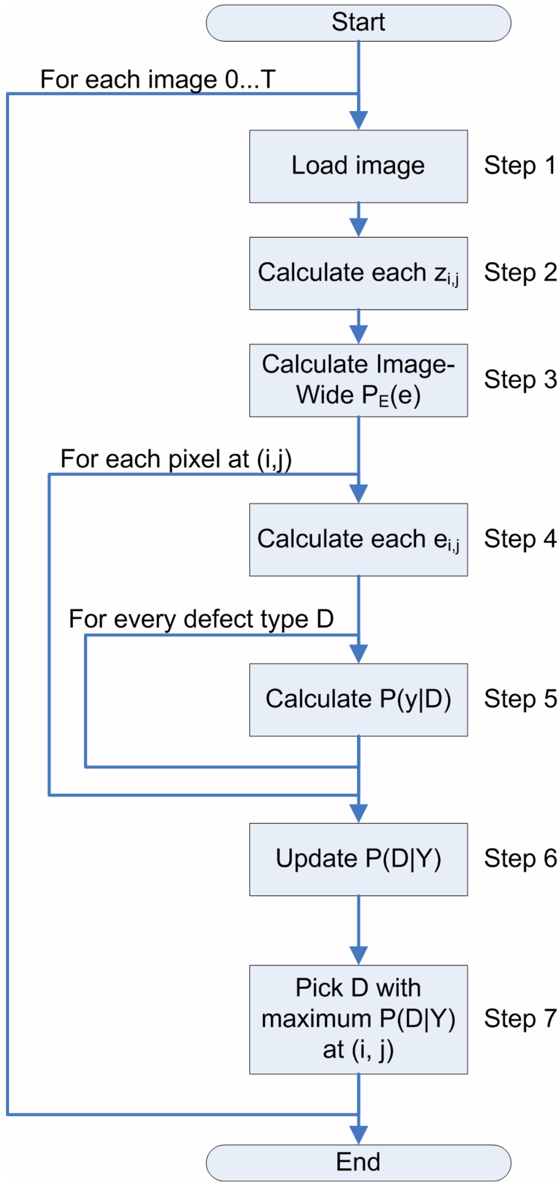

Fig. 17: Flowchart of the defect detection scheme [11]

4.4.3 Extension for Hot Pixels

Leung et al. [14] found out that most observed in-field defects are hot pixels. Therefore they extended or actually modified the above method to distinguish between good and hot pixels only, thus being more accurate at the detection of hot pixels [15]. The transfer function of a hot pixel is:

$$f_{Hot-Pixel}(x) = m \cdot x + b + I_{Hot} T_{Integration}$$

For simplicity only a hot pixel with no additional offset (which is called partially-stuck hot pixel by the authors then) is modelled. The output y of a pixel can then be modelled by:

$$y = x + (T_{exp} \cdot I_{Dark}) = x + \Delta$$

where a good pixel has $\Delta = 0$.

This is an extension of the interpolation method: Like before an interpolation image is generated to estimate the expected pixel value based on its neighbouring pixels' values. Then the interpolation error is calculated using the actual pixel value by $e_{i,j} = z_{i,j} - y_{i,j}$. $p_E(e)$ describes the probability density function (PDF) of this error over the whole image. For a good pixel it should be near 0 if a well performing interpolation scheme is used. If $y_{i,j}$ is the output of a hot pixel then the error should be approximately Δ.

In the following, the subscripts (i, j) will be omitted for simplicity. y denotes the value of a given pixel at (i, j) and y_k denotes the output of this pixel in the k-th image. Now there are only good and hot pixels, thus the Bayes formula for a pixel to be good after the k-th image can be written as:

$$P(Good|y_k) = \frac{P(y_k|Good) \cdot P(Good|y_{k-1})}{P(y_k|Good) \cdot P(Good|y_{k-1}) + P(y_k|Hot) \cdot P(Hot|y_{k-1})}$$

Another consequence is that the likelihood of a pixel being a hot one is simply the compliment of $P(Good|y_k)$:

$$P(Hot|y_k) = 1 - P(Good|y_k)$$

The conditional probabilities $P(y_k|Good)$ and $P(y_k|Hot)$ are again estimated using the interpolation error PDF $p_E(e)$:

$$P(y_k|Good) = p_E(y_k - z_k)$$

$$P(y_k|Hot) = p_E(y_k - (T_{exp} \cdot I_{Dark}) - z_k) = p_E(y_k - \Delta - z_k)$$

The index k is omitted from here on for simplicity. The above equation suggests that the dark current value is a known constant (I_{Dark} and as a consequence Δ is known), but in practice the dark current will change significantly from image to image due to temperature variations and different exposure times additionally influencing Δ. So the computed error $e = y - \Delta - z$ will not be an accurate estimate of the error caused by a hot pixel over several images. Thus the dark current or Δ is modelled as a fluctuating quantity instead, by the procedure described below.

The dark current is measured at calibration time for each pixel at first. Then the first iteration is started with a conservative underestimate of the dark offset denoted by Δ_{min}. With Δ being the dark offset, the estimate of $P(y|Hot)$ is corrected using an upper and lower bound of the dark offset by:

$$P(y|Hot) = P(y|\Delta_{min} \leq \Delta \leq \Delta_{max}) = \int_{\Delta_{min}}^{\Delta_{max}} p_E(y - \Delta - z) \cdot p_\Delta(\Delta) d\Delta$$

As the PDF $p_E(e)$ is just a 512 element vector, the integral can be reduced to the following summation:

$$P(y|Hot) = \sum_{\Delta=\Delta_{min}}^{255} p_E(y-\Delta-z) \cdot \frac{1}{255-\Delta_{min}+1} = \frac{1}{255-\Delta_{min}+1} \cdot [\sum_{x=(y-z)-255}^{(y-z)-\Delta_{min}} p_E(x)]$$

assuming a uniform distribution of dark current values. This summation is the cumulative distribution function (CDF) of the interpolation error $p_E(e)$. If the cumulative sum of $p_E(e)$ is calculated and stored, the summation can further be simplified to ($P_E(e)$ denotes the CDF of $p_E(e)$):

$$P(y|Hot) = \frac{1}{255-\Delta_{min}+1} \cdot [P_E(y-z-\Delta_{min}) - P_E(y-z-255)]$$

Because there are no calibration images available, the value Δ_{min} is set manually. If a dark field calibration image is available, Δ_{min} can be set to:

$$\Delta_{min} = m \cdot ((I_{Dark})_{min} \cdot T_{exp} + b)$$

where $(I_{Dark})_{min}$ and b are derived from the calibration image and $m = \frac{ISO_x}{ISO_{calibrated}}$ is a correction factor for the current image, taken at ISO_x. A problem pointed out by the authors [15] is that the performance of the Bayesian statistics may fluctuate significantly due to the dark current magnitude. But on the other hand due to the accumulative nature of the Bayesian inference method, small changes in $P(y_k|Good)$ and $P(y_k|Hot)$ are reflected in the accumulated statistics. Thus also hot pixel defects in images with low exposure time can be detected [16].

4.4.4 Correction using Local Region Analysis

There are many external factors which influence the performance of the detection algorithm, like the complexity of the image scenes, ISO and exposure settings, the dark current magnitude, additional noise, etc. The main factor limiting the accuracy of the detection algorithm are false detections caused by the scene complexity due to large interpolation errors. Especially image regions containing edges or fine details exhibit more colour variations than usual, which leads to high interpolation errors that could be interpreted as being caused by a hot pixel. Another problem with images taken with a high ISO level is that they are grainier due to the sensor noise, which means that a region with similar colours can still result in large interpolation errors. To improve the performance of the detection algorithm, i.e. reduce false positives, in [14] it is proposed to analyse the image on these factors. Just ignoring images containing fine details would discard much useful information. Thus it is more advantageous to only exclude the image areas containing fine details. This can be done by applying a post procedure, which extracts information about the local region around the pixel of interest. To be more specific, the complexity of a local region can be described by the mean and variance of the region. Thus setting a threshold on these statistical values of the neighbouring region can be used to correct the detection, by simply not taking pixels into account if the variance exceeds a predefined threshold. If this approach is used to track the time of the first appearance of a defective pixel, the authors [17, 14] suggest to ignore the whole image if the variance is above a preset threshold.

5 Ageing Simulation Algorithm

The impact of the defects cannot be measured using two images of the same finger or hand, respectively, taken at different points in time, because then all other external parameters and influences would have to be exactly the same. In addition other influences are also present and one cannot distinguish between the impact of the sensor ageing related pixel defects and the other influences, in particular human ageing. So one would need to have identical data (i.e. fingerprint, finger vein or hand vein images), captured at least at two significantly different points in time, but this is practically impossible. It would not be possible to distinguish between the subject's ageing and the sensor ageing. Therefore the only possible solution to study solely the impact of sensor ageing is to take a given set of fingerprint, finger vein or hand vein images, respectively and artificially age them by simulating the sensor defects.

The simulation only includes hot and stuck pixels. The parameters needed for the simulation should have been found by comparing images of the CASIA2009 and CASIA2013 fingerprint database images and extracting defective pixels caused by sensor defects. Unfortunately it became clear that there were different sensors used in 2009 and 2013 for capturing the images. Thus the pixel defect rate and parameters could not be estimated based on the CASIA the images. Therefore, an empiric formula was used to estimate the defect growth rate based on the sensor parameters, described in section 8.4.1.

5.1 Algorithm

Based on the pixel model in section 3.4.1, the algorithm computes the defect matrices C and D and then applies them to the images. An ideal sensor would have C and D as time-invariant zero matrices. But in practice each sensor starts developing defects at some point in time T_0 with a constant defect growth rate. According to Jessica Fridrich [18], this can be modelled by a Poisson process, which means the number of stuck and hot pixel defects can be calculated as follows:

$$n_s(T) = (T - T_0)\lambda_s$$

$$n_{ps}(T) = (T - T_0)\lambda_{ps}$$

where λ_s and λ_{ps} are the growth rate of the stuck and hot pixel defects[2], respectively. According to research done by Leung et al. [13, 16, 15, 19] and also Albert Theuwissen [3, 4], the defects are being independent from each other and do not develop in clusters, so they can be modelled by a uniform distribution in a 2D sensor array. This means that the defective pixel's position $s_k \in w \times h$ of the k-th defect is obtained from uniformly distributed random variables for the simulation. For a stuck pixel the corresponding value according to the position in the matrix C has to be set and for a hot pixel the value corresponding to the position in the matrix D has to be set. a_s is the maximum amplitude of a stuck pixel and a_{ps} is the maximum amplitude (offset) of a hot pixel (based on 8 Bit grey-scale images). Then the corresponding value in the matrix C and D is obtained by

$$C(s_k) = r_a a_s$$

$$D(s_k) = r_a a_{ps}$$

for a stuck or a partially-stuck pixel at s_k respectively, where $r_a \in \mathbb{R} : [0; 1]$ is a uniformly distributed random number.

[2] Bergmüller et al. used the term partially-stuck pixels instead of hot pixels, thus the subscripts are ps. I use the term hot pixel instead but left the notation as it is proposed in their paper.

According to Leung et al. [15], the dark current is known to fluctuate due to heat variations between each image capture. They modelled the dark current with a fluctuation constant V_{dark} and updated the dark offset according to: $\Delta = I_{dark} \cdot V_{dark} \cdot T_{exp}$. As discussed above, I assume during my simulations that the same sensor is used and all environmental conditions like temperature are kept constant. Thus also the dark current value and so the dark offset are constant. Hot pixels are most observable at long exposure times and the dark offset Δ is a function of the exposure time, but as I assume that all the fingerprint, finger vein and hand vein images, respectively are captured using the same exposure time and thus a constant dark offset value can be used.

The simulation uses images captured at a specific point in time T_0 as basis and adds ageing related sensor defects to each image Y_{T_0}. Of course there might already be some defective pixels due to sensor defects contained in the images taken at T_0, but as only the change over time (increase in the number of defects) is of interest, it does not matter which time frame is observed and especially if there are already some defects present at T_0 or not. The results of the later experiments are relative compared to the baseline performance achieved using the T_0 images, so the defects already contained in the images at T_0 do not influence the results of the experiments. The sensor defects according to the defect matrices C and D at a specific point in time T_i are added to the image Y_{T_0}, which results in the aged image Y_{T_i} then. This image captures the same scene as at T_0 but the sensor defects which occurred during the timespan $T_i - T_0$ are also present in the image. For the evaluation of sensor ageing related pixel defects on finger- and hand-image based biometric recognition systems, a sequence of subsequentially aged images $(Y_{T_i})_{i=1...m}$ is generated by applying the defect matrices C and D to the images taken at T_0. Therefore, at first the sequences of the defect matrices C and D, each representing the state of the sensor at a specific point in time T_i has to be calculated. This sequences are denoted as

$$(D_{T_i})_{i=0...m} \, and \, (C_{T_i})_{i=0...m}$$

The following listing (Listing 1) describes the original sensor ageing algorithm in pseudo code as proposed by Bergmüller et al. [22] to generate the defect matrices and also the sequence of aged images, based on a single source image Y_{T_0} (defect matrices can simply be applied to a set of source images to generate sets of aged images at the specific points in time) and sample points in time $T_0 ... T_m$:

As the number of defects always increases and existing defects do neither disappear nor heal over time, the defect matrices C and D are calculated recursively, i.e. defects are only added and earlier defects are maintained.

The first step at each iteration is to calculate the number of additional defects Δn_s and Δn_{ps} compared to the previous step. Therefore, the locations and the amplitudes of the new defects in the i-th ageing step are determined using random numbers and the new defects are stored in the defect matrices C_{T_i} and D_{T_i}. These are then used to calculate the aged image Y_{T_i}. According to the pixel model, the incident illumination I would have to be used, but as the incident illumination which lead to the image Y_{T_0} is not known, the image Y_{T_0} is used directly instead of I using the argument of relative evaluation again. The output of the algorithm is a sequence of aged images $(Y_{T_i})_{i=1...m}$.

For the present work an extended version of this algorithm with a few additions was used. A minimum distance between two defects can be set to avoid clustering of sensor defects which does not occur in practice (e.g. Leung et al. showed that the minimum distance between two neighbouring defects is between 79 to 340 pixels [13, 15]) but might occur in the simulation. Dudas et al. [11] used the restriction that no two defects may fall in the same 3x3 pixel region while testing their algorithm using Monte Carlo Simulations. This is done using the $checkNeighbours(pos, size)$ function in the pseudo code. In addition the amplitudes or the offsets of the hot pixels, respectively are not drawn from a uniform distribution, but from an exponential distribution, which is denoted as $expd(\mu)$ in the pseudo code. The extended algorithm is shown in the second listing (Listing 2).

Listing 1: Sensor ageing simulation algorithm

```
procedure AgedImageSequence($Y_{T_0}$)
   for i=1...m do
      $\Delta n_s \leftarrow n_s(T_i - T_0) - n_s(T_{i-1} - T_0)$
      $\Delta n_{ps} \leftarrow n_{ps}(T_i - T_0) - n_{ps}(T_{i-1} - T_0)$
      $D_{T_i} = D_{T_{i-1}}$
      $C_{T_i} = C_{T_{i-1}}$
      for k=1...$\Delta n_s$ do
         $r_a \leftarrow$ random in [0;1]
         $s_k \leftarrow$ random in $w \times h$
         $C_{T_i}(s_k) \leftarrow r_a \cdot a_s$
      endfor
      for k=1...$\Delta n_{ps}$ do
         $r_a \leftarrow$ random in [0;1]
         $s_k \leftarrow$ random in $w \times h$
         $D_{T_i}(s_k) \leftarrow r_a \cdot a_{ps}$
      endfor
      $Y_{T_i}(x,y) = \begin{cases} C_{T_i}(x,y) & if\, C_{T_i}(x,y) \neq 0; \\ Y_{T_0}(x,y) + D_{T_i}(x,y) & otherwise. \end{cases}$
   endfor
   return $(Y_{T_i})_{i=1...m}$
end procedure
```

Listing 2: Extended sensor ageing simulation algorithm

```
procedure AgedImageSequence($Y_{T_0}$)
  for i=1...m do
    $\Delta n_s \leftarrow n_s(T_i - T_0) - n_s(T_{i-1} - T_0)$
    $\Delta n_{ps} \leftarrow n_{ps}(T_i - T_0) - n_{ps}(T_{i-1} - T_0)$
    $D_{T_i} = D_{T_{i-1}}$
    $C_{T_i} = C_{T_{i-1}}$
    for k=1...$\Delta n_s$ do
      $r_a \leftarrow$ random in [0;1]
      do
        $s_k \leftarrow$ random in $w \times h$
      while checkNeighbours($s_k$, 3) = 1
      $C_{T_i}(s_k) \leftarrow r_a \cdot a_s$
    endfor
    for k=1...$\Delta n_{ps}$ do
      $r_a \leftarrow$ drawn from expd(0.15)
      do
        $s_k \leftarrow$ random in $w \times h$
      while checkNeighbours($s_k$, 3) = 1
      $D_{T_i}(s_k) \leftarrow r_a \cdot a_{ps}$
    endfor
    $Y_{T_i}(x,y) = \begin{cases} C_{T_i}(x,y) & if\, C_{T_i}(x,y) \neq 0; \\ Y_{T_0}(x,y) + D_{T_i}(x,y) & otherwise. \end{cases}$
  endfor
  return $(Y_{T_i})_{i=1...m}$
end procedure
```

6 Fingerprint Recognition

6.1 Biometric Recognition Systems

The aim of biometric recognition systems is to identify human beings using some kind of biometric trait. A biometric trait is based on a certain distinctive anatomical or behavioural characteristic of an individual. The most common modalities are fingerprint, iris, face, palmprint, finger vein, etc. A biometric trait cannot be shared or misplaced and is therefore ideally suitable for identification and authentication purposes. It also helps in reducing fraud and abuses and improves user convenience, as users do not have to remember multiple complex passwords any more. In addition, only by applying biometrics, negative recognition, i.e. the system checks if the individual is the one it denies to be, is possible. There are some important properties which are either desired or necessary for biometric traits:

- Uniqueness: Is the trait distinctive across users? This means that the trait must be distinctive enough between two different subjects, which is called inter-person variability. This variability should be as high as possible, i.e. the trait has a high discriminative power, to achieve a low number of misclassifications.
- Permanence: Does the trait change over time? There is an intrinsic variability for each biometric trait from one capture to another for the same person, which is called intra-person variability. This is due to ageing processes and different environmental conditions at the point in time of capture. Of course this variability should be as low as possible.
- Stability: qualitative property of a trait with high distinctiveness and high permanence.
- Universality: Does each person have the biometric trait?
- Collectability: Can the trait be measured quantitatively? This is mainly provided by the sensor technology used, but also the processing chain has to guarantee a certain quality. The time that it takes to acquire the sample of the trait is also important and should be as little as possible.
- Performance: Does it meet error rate, throughput? This is mainly influenced by the speed of sensor technology, processing, matching and scalability in terms of the number of users, detection accuracy and resources needed to provide a given accuracy.
- Acceptability: Is the capturing of the trait acceptable to the users? Are the users willing to use the system if they have to? The system should be non-invasive, require only a minimal level of user experience to use it, should be easy to use in general and of course be convenient for the users.
- Vulnerability: Can the trait be easily spoofed, attacked or even circumvented?
- Integration: Can the trait be embedded in the application? The question here is if the trait can be applied in an existing application with some adoptions or if it is not usable for a given application at all. E.g. a fingerprint based authentication system is not the best choice for coal mine workers, because the dirt and abrasion due to their work might make the capturing of a fingerprint image with sufficient quality impossible. Here an iris scanner or a hand- or finger vein based system would be a better choice.

A biometric recognition system operates in two steps. The first step is the enrolment process, where the biometric samples of a subject are captured for the first time and stored in a database under the ID of the subject. Most of the time not the raw input data, e.g. a fingerprint image is stored, instead there is a feature extraction stage, at which certain features, e.g. the minutiae points are extracted, preceded by a quality check and a preprocessing step. Then only the extracted features are stored in the database, which are called a biometric template. The raw input data can be stored in addition to the template, which can optionally be encrypted and compressed, to be able to do an alternative feature extraction at a later point in time. The next step is the authentication procedure, where another sample is captured and used to verify or determine the identity of a given subject. In the verification scenario the captured sample is compared against the template stored inside the database under the ID the subject claims to have, i.e. a 1 to 1 comparison. The same steps like during enrolment are performed, i.e. also a quality check, preprocessing and feature extraction is done prior to calculating the match score using the extracted features and comparing these to the ones stored inside the database. According to this comparison, a matching score is calculated and using a threshold, the system decides whether it is the same subject or not. In contrast to verification, there is also the identification scenario: The current capture of the sample is from an unknown subject and is compared to each sample stored in the database to find the best match, which has a certain likelihood to be from the same subject as the current capture, i.e. a 1 to N comparison is performed. In this scenario it is not sufficient to return the ID of the template, which achieved the highest matching score and say that this is the unknown subject's ID. A likelihood test is also necessary to determine the probability that it is really this subject, depending on the matching scores of the current capture against all other templates inside the database. Especially for the identification scenario a very low EER is essential to get a reliable result.

6.2 Fingerprint Recognition Systems

In fingerprint recognition the used biometric trait are images of the subject's fingers, more precisely the impression of the friction ridges on the inside of the fingertips. Fingerprints meet the requirements for biometric traits as they are highly distinctive and unique, they are supposed not to change severely during the lifetime of a person, are publicly accepted as reliable, their acquisition is easy and convenient for the users and fingerprint sensors can be integrated into most existing authentication systems. The vulnerability can be reduced by performing a fake fingerprint detection using e.g. liveliness checks, which can be done based on physical (e.g. elasticity), visual (e.g. colour) or electrical (e.g. resistance, capacitance) properties of the skin or based on generated signals like pulsation, blood pressure, perspiration, etc. Fingerprint images are captured using different types of scanners, e.g. optical sensors, which work most like a digital camera or sweep sensors, integrated into most modern notebook computers. Also silicon based sensors are common. Today the most widely used sensors are optical ones working based on the FTIR principle. Like all other biometric recognition systems, fingerprint recognition systems operate in two steps, enrolment and authentication. The second step is different for verification and identification. The schematic procedures for enrolment, verification and identification can be seen in figures 18, 19 and 20. As quality metric the NFIQ (Nist Fingerprint Image Quality) is the de facto standard for evaluating the quality of fingerprint images. To achieve a higher robustness, not only the image of one finger of each person is captured during enrolment, but at least two fingers are used. The subject may hurt one finger and would then not be able to authenticate himself if only one finger is enrolled. In addition in some applications, a re-enrolment may be possible to store an updated version of a subject's fingerprint, whereas in others this might be prohibited, like in high-security applications to avoid forging attacks.

6.2.1 Fingerprint Sensors

Here the different types of fingerprint sensors which are available should be discussed briefly with a focus on optical sensors using a CCD/CMOS image sensor as only these sensors are affected by the

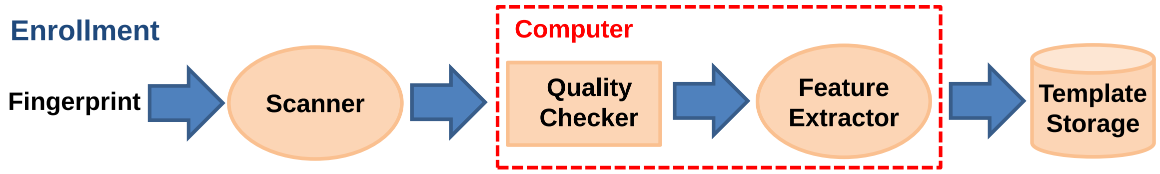

Fig. 18: Enrolment step

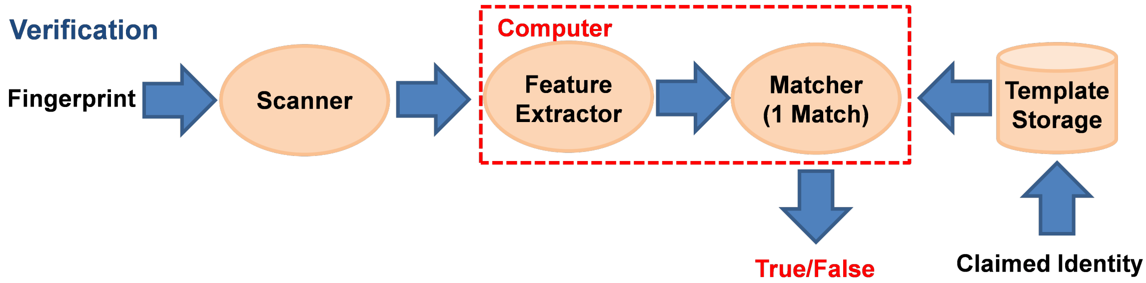

Fig. 19: Verification process

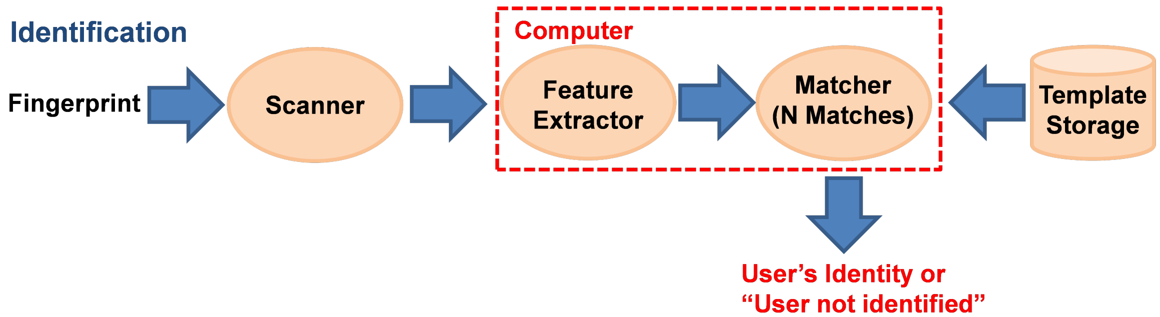

Fig. 20: Identification process

kind of sensor ageing related pixel defects as described in section 3.4. Of course also thermal and other sensors are affected by ageing related effects but to the best of my knowledge the impact on the sensor output has not yet been investigated and there is no defect model available. Consequently the focus of this thesis is on biometric sensors using a CCD or CMOS image sensor.

Off-line acquisition In forensics, so called latent fingerprints are captured off-line. If a finger touches an object, there remains a moisture or grease film on it, which renders the ridge and valley structure. This film can be made visibly by different methods, e.g. using powder or chemical treatment.

Also for authentication systems the so called ink technique is used to capture fingerprints off-line. Therefore the finger is put onto an ink pad and then pressed against a cardboard. Here the whole fingerprint, from nail to nail can be captured which stores more information than by using an FTIR sensor. The quality of the fingerprint not only depends on the uniform dispensation of the ink, but also on the condition of the finger itself (fat, dirt, sweat, moisture). The print on the cardboard is then converted into a digital image using a scanner or camera.

Optical - FTIR and Direct FTIR, which stands for frustrated total internal reflection, is the most common used fingerprint sensor technology. The finger is placed on top of a glass prism, the ridges directly touch the prism and the valleys form a small gap, where there is air closed up between the skin and the prism surface. One side of the prism is illuminated, the light is being absorbed at the ridges and reflected at the valleys, thus the valleys appear bright and the ridges appear dark. The reflected light passes out of the prism on the other side, where it is focused on a CCD image sensor through a lens. The optical distortion, which is introduced by the prism, has to be corrected optically or by software. The size of the prism can be reduced by using a sheet prism, but this reduces the quality.

The Direct sensor technology uses an optical fibre with a CCD sensor directly applied to it, which replaces the prism and the lens. Residual light originating from the finger is being transferred through the optical fibre and captured without the need of an external light source. The large CCD sensor area leads to high costs, which is the main disadvantage.

Electro-Optical These sensors consist of two layers. The first one is a light emitting polymer, which light intensity depends on the electric potential applied on the side of it. The ridges touch the surface whereas the valleys not, so they have a different potential and thus the emitted light has a different intensity. The second layer consists of photodiodes, arranged in an array, which capture the emitted light and transform it into an electric signal, which is then converted to the digital image. These sensors can be built in very small sizes, but the quality is inferior compared to FTIR sensors.

Silicon - Capacitive Silicon based sensors are also called solid-state sensors and consist of an array of pixels where each pixel is a sensor on its own. The finger directly touches the silicon surface. They only differ in the way the physical information (ridge and valley structure) is translated into electrical signals.

A capacitive sensor consists of an array of micro capacitors, in particular only one capacitor plate. The second plate is the finger itself. If the finger touches the surface, this results in small electrostatic charges and according to the laws of the electrostatic field, the magnitudes of the resulting charges are dependent on the distance between the two plates, i.e. the distance between the sensor's surface and the skin surface of the finger. Thus the ridges and valleys cause charges with a different magnitude which can be measured. The surface has to be protected using a protection layer and it has to be resistant against electrostatic discharges. These sensors cannot be tricked by a photograph of the finger.

Silicon - Thermal Each individual sensor is made of a pyro-electrical material, which generates an electrical current dependent on temperature differences. The sensor is kept at a certain temperature, higher than room temperature, to provide a high enough temperature difference to the skin temperature. As the ridges directly touch the sensor surface and the valleys not (the air in between isolates it thermally), the sensor surface is cooled down where the ridges are. Due to the temperature compensation the image is only visible a short amount of time, even if the finger still stays on the sensor. Thus a dynamic capturing method like sweeping should be used. These sensors are more robust against discharges and the protective layer can be thicker.

Silicon - Piezoelectrical Each pixel contained on the sensor is pressure-sensitive and the electric output signal depends on the pressure applied to each pixel. The pressure at the ridges is higher than at the valleys, which causes a different output signal. The currently available materials are not sensitive enough, the problem with the protective layer is not yet solved and the output is a binary image.

Silicon - Electric Field The sensor generates an electric field, which is distorted by the structure of the skin surface. These distortions can be measured and converted to an output signal. This type of sensor is also vulnerable to electrostatic discharges.

Ultrasonic Ultrasonic sensors work with acoustic signals in the ultrasonic range, which are reflected at the ridge and valley structure of the finger. The reflected sound waves are captured by a sensor and based on the different points in time when they are captured, the distances that the signals have travelled can be determined. Sound waves reflected at ridges have shorter distances than the ones reflected at valleys. The quality of the output images is quite good and is it not influenced by dust, dirt, wearing gloves etc. The main disadvantages are the size of such a sensor, the high costs and the long acquisition time of a few seconds.

Touch and Sweep Touch means the finger is simply placed on top of the sensor surface. No training of the user is necessary but the sensor surface may get polluted fast. This reduces the image quality, moreover the finger can be placed rotated or misaligned and a latent fingerprint resides on the sensor after the capture.

Sweep means that the finger is moved over the sensor in a vertical sweeping motion. The sensor has to be as wide as the finger and should be at least several pixels high to enable a robust reconstruction of the image from the slices. The production costs are lower, the sensor does not get polluted so easily as it is constantly cleaned while sweeping and the finger cannot be placed rotated. However the user needs some training to perform an even sweeping motion with the right speed, the sensor has to be fast enough to capture the image slices while sweeping and the image has to be reconstructed from the slices, which may introduce additional errors.

6.2.2 Fingerprint Image Examples captured by Different Sensors

Figure 21 shows some fingerprint images acquired with different types of fingerprint sensors. HiScan, MSO350, UareU4000 and Verifier 300 LC 2.0 are FTIR sensors. FingerChip AT77C101B is a thermal sweep sensor. AES 4000 is a representative of an electric field sensor and TouchChip TCS1 is a capacitive sensor.

(a) HiScan (b) MSO350 (c) FingerChip AT77C101B (d) Verifier 300 LC 2.0 (e) AES 4000 (f) UareU4000 (g) TouchChip TCS1

Fig. 21: Example images captured by different fingerprint scanners [26]

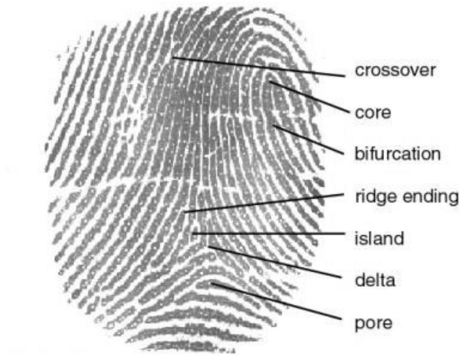

Fig. 22: Minutiae types [27]

6.3 Fingerprint Anatomy

The skin on the inside of human feet and hands is covered with concentric raised patterns, forming ridges and valleys. These ridges provide friction making it easier for humans to hold and grasp objects. The pattern shown in figure 22 which is formed by these ridges is called a fingerprint. A fingerprint consists of mainly parallel flowing lines called ridge lines, forming a pattern called ridge pattern. These lines are not continuous or straight, instead they are broken, forked or change direction. Points on which the ridges change, fork or end are called minutiae points. These minutiae points, which are extracted and used as features by the most common matchers, are determined by the termination or bifurcations of the ridge lines and can be classified in:

1. Ridge Ending: end of a ridge
2. Ridge Bifurcation: a single ridge divides into two ridges
3. Independent Ridge: the ridge begins and ends again already after a short distance
4. Island: a small ridge inside a short ridge
5. Spur: a bifurcation, at which a short ridge branches out of a longer one
6. Crossover: a short ridge running between two parallel ridges

7. Delta: ridges meeting in a Y-shape
8. Core: a U-turn of a ridge

The most commonly used types are ridge endings and ridge bifurcations. There are other local structures produced by the ridge lines, the most common ones are a whorl, a loop and an arch, which can be used for a coarse classification of fingerprint images and therefore as a first matching step. These singularities are most commonly found by a method based on the Poincaré index.

In addition, there are more local features, which are only visible at a very high resolution of the scanner, e.g. at 1000 dpi, which include sweat pores, creases and incipient ridges. These features can also be used to perform an even more accurate matching but require a high quality and high resolution input image.

6.4 Feature Extraction and Matching

For the traditional minutiae based approach the feature extraction stage begins with a segmentation of the input image to separate the foreground (fingerprint area) from the background. This is followed by deriving an orientation, a fingerprint shape and a frequency image, respectively, from the segmented fingerprint image. Based on the orientation image the singularities, e.g. deltas and whorls are found. The three images are then used to derive a ridge pattern, out of which the different minutiae points are extracted. This can also be done quite simple by first binarising the input image, then performing a thinning operation on the binarised image and the minutiae points can then be detected using a simple image scan (e.g. filtering the image with predefined filters). Most feature extractors use at least the minutiae type, direction and position.

Based on this simple method of minutiae extraction there are some intrinsic problems. First of all a certain amount of information may be lost during the binarisation step. Moreover binarisation techniques provide unsatisfactory results in low-quality image regions and thinning may introduce additional spurious minutiae. Thus other types of feature extraction methods have been developed, especially for low quality images, which do not use minutiae points, but some other features like local ridge features, including local ridge orientation and frequency and also correlation based approaches, which use the global ridge and furrow structure.

Different types of fingerprint recognition schemes react differently to quality degradations, thus fundamentally different types of fingerprint feature extraction and matching schemes are considered to be able to study the influence of sensor ageing related defects on each one of them. After the image is captured, different features are extracted from the image based on the discriminative characteristics a fingerprint contains. The different approaches used during the evaluation are described below.

6.4.1 Correlation Based Matcher

In contrast to the minutiae based matchers, a fingerprint image in its entirety is used at the feature extraction stage, i.e. the global ridge and furrow structure is decisive. The fingerprint images are correlated using different translational and rotational alignments. During the experiments done in this thesis, a custom implementation of the phase only correlation (POC) matcher [28] done by Michael Pober was used. This POC matcher works as follows: At first the DFT (discrete Fourier transform) of each of the two images, which should be matched against each other, is calculated. Afterwards the normalised cross spectrum (or cross-phase spectrum) is calculated in the Fourier domain. Finally the inverse DFT of the normalised cross spectrum is computed, which is the actual POC then. POC has some advantageous properties like shift invariance, brightness invariance and immunity to noise, which makes template alignment easier. The rotational alignment is done by computing the POC for rotated versions of the original fingerprint images in a range of $\pm 20°$ with a step size of $1°$. The rotated version with the highest POC score is used for matching. The POC peak is used to perform displacement alignment and subsequently probe and gallery images are cropped to the common area of intersection, as non-overlapping regions would lead to uncorrelated noise in the POC function. For the final score a band limited version of the POC is computed. Band limiting is done in a way trying to limit the frequency spectrum to only those regions in the fingerprint image that are related to actual fingerprint data, especially the ones originating from the ridge and furrow structure and excluding high frequency areas with interfering components. The last step is the calculation of the final matching score which is computed by summing up the P highest peaks. In the custom implementation $P = 1$, i.e. only the highest peak is used. Fingerprint enhancement as suggested in [29] is done before feature extraction to improve the matching results.

6.4.2 Ridge Feature Based Matcher

This type of matcher uses the overall ridge and furrow structure not in a global manner like the correlation based one but in a localised manner. Features like local ridge orientation and local ridge frequency are extracted and used to represent the individual fingerprint. During this thesis, as a representative of this category, a custom implementation of the fingercode (FC) approach [30] done by Michael Pober was used, which is based on Gabor filters. Again with this custom implementation the fingerprint enhancement strategy suggested in [29] is used.

Prior to the actual feature extraction, normalisation is applied to the input images by pixel-wise adjusting the grey-levels to get an image with a specific mean and variance. A least squares estimate of the local ridge orientation, computed in 16 x16 pixel blocks is then calculated using a Marr-Hildreth operator. This is followed by applying a low-pass operator to smoothen the output, which forms the so called orientation image. This orientation image in combination with the normalised fingerprint image is then used to create the frequency image, containing the local ridge frequency. Then the x-signature is calculated per block using an oriented window by projecting the respective grey levels of the normalised fingerprint image onto the length of the window, which is placed in a direction orthogonal to the local ridge orientation of each block. Using the reciprocal of the average distance between the peaks in x-signature, the frequency is determined. In addition interpolation of invalid blocks, i.e. blocks where the x-signature does not follow a discrete sinusoidal-shape wave, is done using a discrete Gaussian kernel.

Feature extraction is based on a bank of eight separate Gabor filters, each one oriented at a different, constant angle. The image is convolved with each of the eight filters, resulting in eight distinct filtered images, each of them containing the response of the particular Gabor filter applied to the ridge and furrow structure of the fingerprint. For each image the standard deviation at every pixel position in a 16 x 16 block, centred at the pixel under investigation, is computed. The union over all eight standard deviation images is denoted as Standard Deviation Map. During the enrolment step the image is subsampled by a factor of 16 to generate ridge feature images, which union forms the Ridge Feature Map. Differently displaced ridge feature images are correlated in the Fourier domain and by identifying the image with maximal correlation, compensation for shift and translational alignment is done. Rotational alignment is performed by simply storing a set of ridge feature maps calculated from rotated versions of the fingerprint images in a range of $\pm 20°$ with $1°$ step width and then again using the one achieving the maximum correlation value. Finally the matching score is computed by calculating the Euclidean distance between the aligned ridge feature map entries.

6.4.3 Minutiae Based Matcher

Each minutia of the fingerprint is determined, represented at least by its type, location and direction and stored in a list. The matching process then simply tries to find an optimal alignment between two sets (lists) of minutiae points of the two fingerprints to be matched, resulting in a maximum number of pairings (including some kind of distance metric) between the minutiae from the one set to compatible ones from the other set.

NFIS2 (mindtct and bozorth3) The first representative of a minutiae based matcher used in this work is mindtct (feature extraction) and bozorth3 (matching), both contained in the "NIST Biometric Image Software (NBIS)" package, available at `http://fingerprint.nist.gov/NBIS/`. mindtct locates and extracts minutiae points contained in the input fingerprint image and assigns type, location, orientation and quality to each point. The process can be divided into 6 steps:

1. Generation of image quality map
2. Binarisation
3. Minutiae detection
4. Removal of false minutiae including lakes, holes, islands, hooks, overlaps, side minutiae, minutiae in low quality image regions, minutiae that are too wide or too narrow (pores)
5. Counting of ridges between each minutiae point and its nearest neighbours
6. Minutiae quality assessment

The minutiae detection is unreliable in low quality image regions. Therefore, the first step is to determine image areas that have low quality using characteristics like low contrast, high curvature and incoherent ridge flow. These characteristics represent unstable image areas and are used to assign a quality level to each image area in order to generate an image quality map. The image is divided into non-overlapping blocks and to each block a quality level from 1 to 5 is assigned. This is followed by a binarisation step. Afterwards the minutiae detector scans the binarised image, trying to identify local pixel patterns that indicate splitting or ending of a ridge. The candidate minutiae points are then detected by comparing the patterns to a set of minutiae patterns. In the next step, false minutiae points are removed and the remaining ones are considered as correct minutiae of the fingerprint. Apart from the minutiae information itself, mindtct also extracts additional information, in particular the ridge count between each minutiae point and its nearest neighbours. In the last step, a quality score is assigned to each of the remaining minutiae points, because some false minutiae may still remain in the list after the removal stage. The quality score is based on two measures. The first one is taken directly from the quality map of the image at the location of the minutia point and the second one is based on pixel intensity statistics of the minutia point's neighbourhood, i.e. mean and standard deviation.

bozorth3 is the matching algorithm which computes a match score between all the minutiae points from two fingerprints. Although mindtct extracts type, location, orientation and quality of each minutia point, bozorth3 uses only the location and orientation information of the minutiae points to calculate a match score. It is rotation and translation invariant. The matching score is based on computing appropriate scores for each pair of minutiae points of the two images, based on the distance between the two minutiae and the angle between each minutia's orientation and the intervening line between both minutiae.

VeriFinger (Neurotechnology) The second representative is a commercial one, developed by Neurotechnology and called VeriFinger, which is available as trial version of the VeriFinger SDK, version 7.0 at `http://www.neurotechnology.com/download.html#verifinger_sdk_trial`. According to Neurotechnology's website[3]: *"The VeriFinger algorithm follows the commonly accepted fingerprint*

[3] `http://www.neurotechnology.com/verifinger-technology.html`

identification scheme, which uses a set of specific fingerprint points (minutiae) along with a number of proprietary algorithmic solutions that enhance system performance and reliability." VeriFinger is tolerant to fingerprint deformations and able to match flat against rolled fingerprints and vice versa. It is robust against translation, rotations, and deformation. In addition stuck ridges, ridge ruptures and noise are eliminated prior to minutiae extraction to get more reliable results even from poor quality fingerprints. Even if the template and the query fingerprint image have 5 – 7 matching minutiae only an identification should still be possible. It uses a feature generalization mode, i.e. features are extracted from several fingerprint images of the same finger, analysed and then combined into a single generalized feature collection, which is stored in the database. This makes the enrolled features more reliable and increases the recognition quality.

MCC - Minutia Cylinder Code The Minutia Cylinder Code approach is no combined feature extraction and matching technique. Instead it is a new local matching approach for previously extracted minutiae points. In contrast to global minutiae matching algorithms, which are computationally demanding and not robust against non-linear fingerprint distortions, local ones use attributes that are invariant to global transformations, e.g. rotation, translation, scaling and aim to be less computationally demanding. But as only local minutiae arrangements are used for matching, the information contained in global spatial relationships is not available and therefore there is less discriminative information available for matching the fingerprints. This might lead to a lower matching performance compared to the usual global approaches. This approach could be used in a hybrid two-stage strategy. At first the local matching is used to reduce the number of candidate matches and then a global matching approach is used to make the final decision.

According to the Cappelli et al. [31] the main advantages of their approach are:

- *it is a fixed-radius approach and therefore it tolerates missing and spurious minutiae better than nearest neighbor-based approaches;*
- *unlike traditional fixed-radius techniques, it relies on a fixed-length invariant coding for each minutia and this makes the computation of local structure similarities very simple;*
- *border problems are gracefully managed without extra burden in the coding and matching stages;*
- *local distortion and small feature extraction errors are tolerated thanks to the adoption of smoothed functions (i.e., error tolerant) in the coding stage;*
- *it effectively deals with noisy fingerprint regions where minutiae extraction algorithms tend to place numerous spurious minutiae (close to each other); this is made possible by the saturation effect produced by a limiting function;*
- *the bit-oriented coding makes cylinder matching extremely simple and fast, reducing it to a sequence of bit-wise operations (e.g., AND, XOR) that can be efficiently implemented even on very simple CPUs.*

MCC represents each minutia using a local structure, which encodes the spatial and directional relationships between the minutia and its neighbourhood. This structure can be represented as a cylinder with its base and height depending on the spatial and directional information.

Matching is done using a local similarity measure, based on a vector correlation measure. A bit-based version of this measure can be implemented and used on architectures where floating point operations are not available or rather slow to speed up the matching process.

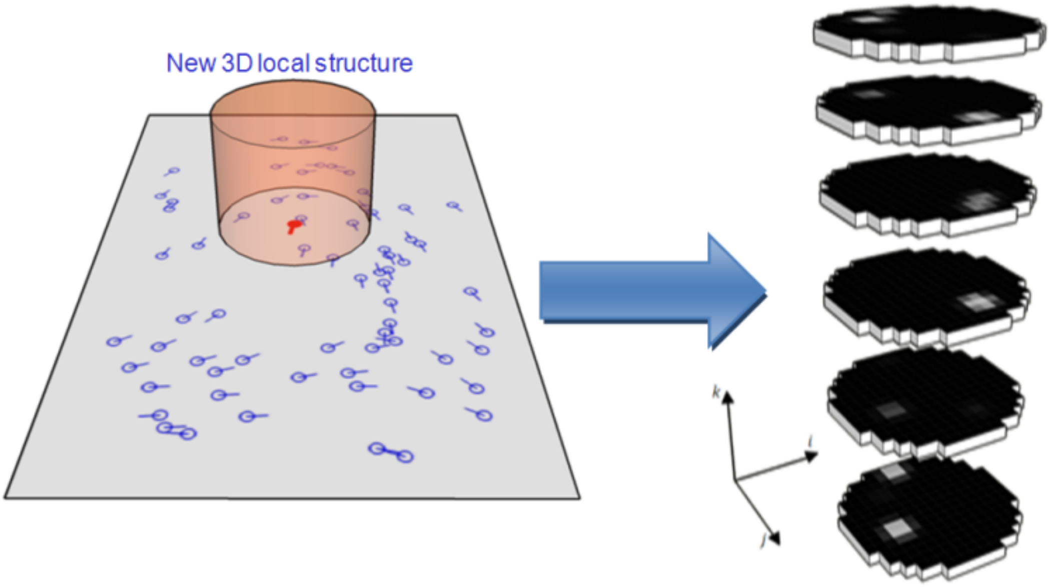

Fig. 23: Local 3D structure representation [31]

$$\gamma(a,b) = 1 - \frac{\|\mathbf{c}_{a|b} - \mathbf{c}_{b|a}\|}{\|\mathbf{c}_{a|b}\| + \|\mathbf{c}_{b|a}\|} \tag{1}$$

$$\gamma_{Bit}(a,b) = 1 - \frac{\|\mathbf{c}_{a|b}\ \mathrm{XOR}\ \mathbf{c}_{b|a}\|}{\|\mathbf{c}_{a|b}\| + \|\mathbf{c}_{b|a}\|} \tag{2}$$

Fig. 24: Matching two cylinders [31]

6.5 Factors Influencing Recognition Accuracy

In the identification (not verification), thus in an 1:n and not an 1:1 matching scenario, a high matching accuracy, i.e. a low equal error rate, is necessary to successfully identify an individual. If everything would be perfect, then either two fingerprints would match exactly or they would not match at all. But unfortunately even though the ridge pattern does not change significantly over time, fingerprint images can vary a lot due to different skin conditions, dirt and abrasion. This can go as far as to have two fingerprints of different fingers look like they are taken from the same finger, e.g. due to creases or worn fingerprints, which makes feature extraction very challenging. Especially in forensic applications, where most of the time latent fingerprints are used, there are many factors that influence the quality of the fingerprint and therefore the recognition accuracy. Moreover, the higher the number of templates, the higher is the likelihood of two fingerprints looking like the same if they are not. So it is vital to know which feature extraction/matching scheme is robust against different types of quality degrading factors, to be able to choose the right one for the given purpose.

6.5.1 Acquisition Area

The most important parameter is the acquisition area. The larger it is, the higher is the accuracy, i.e. the lower the EER. An acquisition area of 240 mm^2 should be the minimal area used. For good results an acquisition area of 350 mm^2 or higher is recommended. Using an acquisition area of only 210 mm^2, the average performance drop in terms of EER is about 70%.

6.5.2 Displacement and Rotation

If there is a high displacement or a high degree of rotation, only a small overlap between the template and the current sample exists, which makes matching difficult even if the input image has a high quality. The problem gets worse for small area sensors. Another problem is motion blur, resulting from small movements of the finger while capturing the image, which can be reduced using short exposure times.

6.5.3 Non-Linear Distortion

The sensor performs a mapping of the 3D structure of the finger onto the 2D surface of the sensor, which results in a non-linear distortion. This distortion is not always the same due to the skin plasticity, which could result in images of two different fingers looking very similar.

6.5.4 Pressure

Especially uneven and non-uniform pressure applied to the finger while put on the sensor surface causes problems, as it further distorts the image in a non-linear way. Some of these effects can be corrected using preprocessing, trying to model the skin distortion at first and then filtering the effects according to the modelled distortion.

6.5.5 Skin Conditions

The condition of the finger's skin, which is placed on the fingerprint sensor may also influence the quality of the captured fingerprint image and therefore the recognition accuracy. Especially dust, cuts and dirt on the finger lead to disruptions in the ridge and valley lines, causing the extraction of minutia points to be more difficult. Dryness of the skin, sweat, grease and high air humidity may lead to a poor overall image quality, depending on the sensor type, which might make the extraction of minutiae points nearly impossible. Thus other feature extraction and matching schemes, e.g. correlation- and ridge-feature based matchers have to be used.

6.5.6 Sensor Noise

Sensor noise may be caused due to a high ISO level, which results in shot noise in the output image. A high ISO level is used to achieve short exposure times in low light conditions, helping to reduce motion blur. This is a problem especially for sensors having a very small pixel size, as each single pixel may only be collecting a few thousand photons. The output of the pixel is then affected by random statistical fluctuations in the photon density, which leads to noise in the output image. Another type of sensor noise is the quantization noise, which is caused by the quantization of a single pixel's output. Additional noise may be introduced by the sensors electronics, errors in signal transmission and errors on the storage medium. Sensor noise is typically also dependent on environmental conditions, especially on the temperature.

Another type of sensor noise are pixel defects due to sensor ageing. These defective pixels further degrade the image quality and their influence on the recognition accuracy of biometric recognition systems is the main topic of this thesis.

6.5.7 Fingerprint Image Enhancement

Sensor noise can be reduced using a denoising filter. If a sensor's characteristics and its noise pattern are known, this can be used to reduce the sensor noise to a great extent by using deconvolution. Displacement and rotation can be corrected by trying to align the images or by computing the matching score for different translations and rotations. The best approach for further enhancement of fingerprint images is using contextual filters, e.g. Gabor Filters, trying to extract exclusively the ridge and valley structure. Despite all enhancement techniques, today humans still outperform automatic fingerprint feature extraction algorithms in extracting features from very low-quality input fingerprint images.

7 Finger and Hand Vein Recognition

For authentication and identification not only fingerprints can be used but also finger and hand vein[4] images. Basically these images capture the pattern of the blood vessels inside the fingers and hands of a human.

Using the vein patterns has several advantages over the well established fingerprints. First of all the veins are underneath the skin and only visible in near infrared light, thus the vein pattern is resistant to forgery. Moreover, liveness detection is easily possible. The most important advantage is that the vein patterns are insensitive to abrasion and finger surface conditions, like moisture, dirt, dryness, cuts and so on.

On the other hand there are also some disadvantages. First of all it is not completely clear if the vein patterns exhibit sufficiently distinctive features to reliably perform biometric identification among large user groups. The data sets which are currently public available, are limited in their size and therefore, this issue cannot be clearly answered by now. The second disadvantage are the capturing devices needed for taking images of the vein pattern. It is necessary to illuminate the human hand or fingers with infrared light to make the inside veins visible. There are two ways of doing this, either by reflected light or using transillumination. Transillumination is able to make also vein lines further underneath the skin visible, leading to more distinctive output images. But scanners using transillumination are rather big compared to fingerprint sensors. An example can be seen in figures 25 and 26. The vein structure is influenced by temperature, physical activity and certain injuries and diseases. Although the impact and effects of these influences on the recognition performance has not yet been studied it is known that especially the width of the veins changes due to temperature changes and physical activity. Thus a suitable feature extraction technique should not depend on the width of the veins.

Nevertheless, there are already some uses of finger vein scanners nowadays, e.g. the UK Barclays Bank uses finger vein scanners to approve transactions in their local branches[5] and now introduced finger vein scanner for their customers to approve online banking transactions as well[6].

7.1 Finger and Hand Vein Scanners

7.1.1 Finger Vein Scanner from Veldhuis et al.

The UTFVP data set provided by R.N.J Veldhuis [33] was captured using a custom designed transillumination device. Finger and hand veins are underneath the skin and not apparent in visible light, but they are visible in infrared light. All capturing devices utilize the fact that blood has a higher absorbency than surrounding tissue in the near infrared spectrum, so the veins appear as dark lines, which is due to the haemoglobin in the blood. Images are thus captured using near infrared light, typically with a wavelength between 780 and 930 nm. Veldhuis et al. used infrared LEDs with a wavelength of 850 nm.

Their finger vein scanner consists of an infrared sensitive camera with an additional filter for visible light, a mirror (to minimize the height of the device) and eight individual infrared LEDs, each with its own control loop to achieve uniform light intensity along the finger. Figure 26 shows a schematic and an image of the capturing device.

[4] Actually it are not only the veins that are used for authentication but also the arteries so the more correct term would be hand and finger vascular pattern.

[5] `http://www.theguardian.com/business/2014/sep/05/barclays-introduce-fingervein-id-readers`

[6] `http://www.newsroom.barclays.com/Press-releases/Barclays-first-in-UK-to-launch-new-Biometric-Reader-for-customers-bab.aspx`

Fig. 25: Barclays Bank finger vein scanner [32]

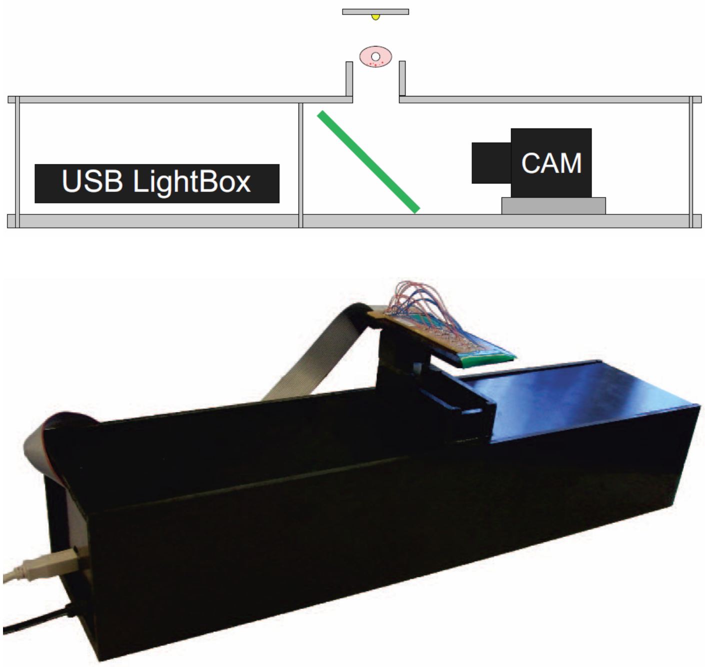

Fig. 26: Image capturing device for finger vein images [33]

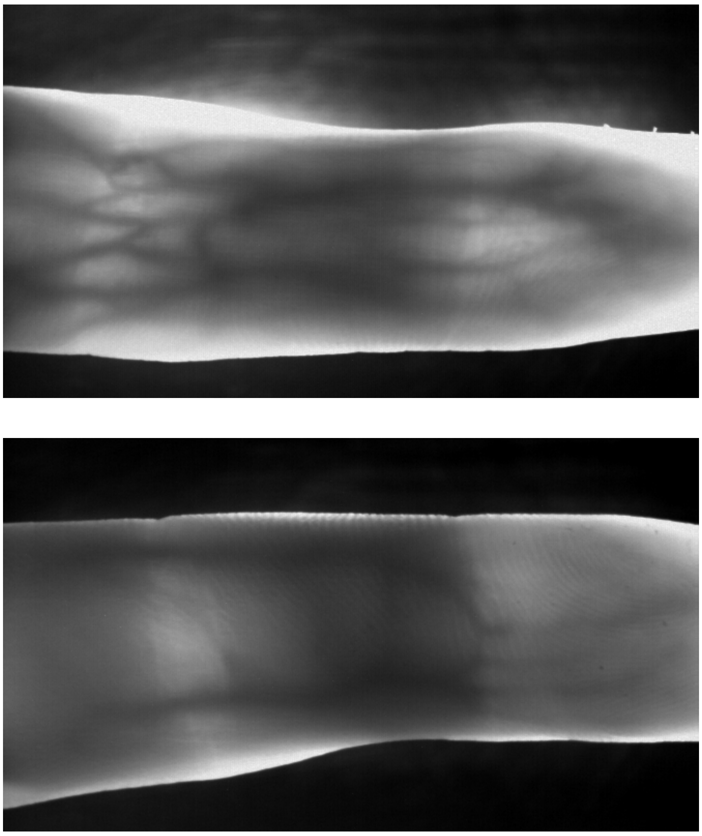

Fig. 27: Example finger vein images [33]

7.1.2 University of Salzburg Hand Vein Scanner

The hand vein image database used during this work was collected at the University of Salzburg using a custom made prototype hand vein scanner. This device basically consists of an infrared light source underneath a glass plate with pegs for alignment, where the hand is placed, for transillumination and another infrared light source on top of the device for reflected light. A modified DSLR is used to capture the images. The infrared light filter of the DSLR has been removed and an additional filter to cut off visible light has been added. All the components are fitted in a wooden box to block off the ambient light, which also has some infrared component and might infer with the controlled light sources.

7.2 Image Preprocessing

One of the main problems regarding finger and also hand vein images is that although a capturing device with an individually controllable light source is used, the images have a rather low contrast and low quality in general. It highly depends on the subject how well the vein structure is visible in the captured images. Figure 27 shows two example finger vein images, a rather good one at the top and a bad one at the bottom, at which the vein structure is only barely visible.

To achieve a higher recognition accuracy, the images have to be preprocessed prior to the feature extraction stage, with the aim to make the vein structure more visible and make it easier for the feature extraction methods to accurately extract it.

Several preprocessing methods were used, starting from the detection of the finger region and borders and aligning the image based on the finger edges to compensate rotations. The most basic method used for contrast enhancement is an adaptive histogram equalisation (CLAHE), which is widely used in the literature. In addition some methods specifically designed to enhance the vein pattern are used, which are High Frequency Emphasis Filtering and a Circular Gabor Filter based technique. Furthermore, sharpening using Unsharp Masking and denoising using a median filter and an adaptive Wiener filter were tested.

7.2.1 Detecting Finger Region (LeeRegion [1])

LeeRegion is a quite simple method for detecting the finger region and setting all pixels outside it to black, which is favourable for further preprocessing (histogram equalisation). It uses the fact that the finger region is brighter than the background and determines the finger boundaries using a simple 20 x 4 mask, containing 2 rows of 1 followed by 2 rows of -1 for the upper boundary and a horizontally mirrored one for the lower boundary. The position at which the masking value is maximal is determined as the finger edge (boundary).

7.2.2 Finger Position Normalisation [2]

Due to slight variations in the positioning of the finger on the capturing device, the orientation of the finger does not always have to be the same. Therefore, this method tries to align the finger to the centre of the image, compensating rotations and translations. It uses the finger edges detected by LeeRegion and attempts to fit a straight line between these detected edges. The parameters of this line are then used to perform an affine transformation, which aligns the finger.

7.2.3 CLAHE (Contrast Limited Adaptive Histogram Equalization)

CLAHE is suggested by most authors as a preprocessing step for finger vein images due to their low contrast. CLAHE is a local contrast enhancement technique, similar to histogram equalization. Standard histogram equalization modifies global image contrast to achieve a nearly uniform histogram. It does not work well with images containing bright areas with high detail and low contrast background. In this case details are lost. CLAHE divides the image into small blocks, each block is then histogram equalized separately. Contrast limiting is applied to avoid amplification of noise: any histogram bin over a defined limit is clipped and those pixels are distributed uniformly to other bins. As a final, step bilinear interpolation is used to reduce border artefacts at the block boundaries.

7.2.4 High Frequency Emphasis Filtering

High Frequency Emphasis Filtering (**HFE**) is applied in the Fourier domain. It was first proposed by Zhao et al. [34]. They used it for hand vein image enhancement. At first the discrete Fourier transform $F(u,v)$ of the image $f(x,y)$ is calculated. The transform is then centred (image is multiplied by $(-1)^{x+y}$) and the high frequency emphasis filter:

$$G(u,v) = H_{he}(u,v) \cdot F(u,v)$$

$$H_{he}(u,v) = a + b \cdot H_{hp}(u,v)$$

where

$$H_{hp}(u,v) = \frac{1}{1 + (\frac{D_0}{D(u,v)})^{2n}}$$

is a Butterworth high-pass filter of order n, is applied. a , b , D_0 and n are the parameters of the filter. Afterwards the inverse Fourier transform of the enhanced image $G(u, v)$ is computed and its real part is used as result. Finally the resulting image is multiplied by $(-1)^{x+y}$ to get the enhanced image $g(x, y)$. The authors used histogram equalization on the filtered image afterwards to further enhance the image contrast.

7.2.5 Circular Gabor Filter

The Circular Gabor Filter (**CGF**) was proposed by Zhang et al. [35] in combination with grey level grouping for contrast enhancement. Here only the Circular Gabor Filter is used and CLAHE is performed afterwards instead of grey level grouping.

Circular Gabor Filters are rotation invariant and can achieve optimal joint localisation in both, spatial and frequency domain. Here an even-symmetric circular Gabor filter is used:

$$G_C(x, y) = g(x, y) \cdot cos(2\pi f_c \cdot \sqrt{x^2 + y^2})$$

where

$$g(x, y) = \frac{1}{2\pi\sigma^2} \cdot e^{-\frac{x^2+y^2}{2\sigma^2}}$$

is a 2-D isotropic Gaussian envelope. The Relation between σ and f_c is:

$$\sigma \cdot f_c = \frac{1}{\pi} \cdot \sqrt{\frac{ln(2)}{2}} \cdot \frac{2^{\Delta F} + 1}{2^{\Delta F} - 1}$$

σ and ΔF are the parameters of the filter. The enhanced image is then computed using convolution:

$$F(x, y) = G_C(x, y) \star I(x, y)$$

where $\star$ denotes the 2-D convolution operator.

7.2.6 Further Preprocessing

A median filter (5 x 5) and an adaptive Wiener filter (7 x 7) for denoising were tested but did not improve the results except for LBP were the EER could be improved. As the implementations provided by B.T. Ton [36] resized the images by a factor of 2, the images were also resized to half of its original size in order to comply with the test procedure and to lower the runtimes.

7.2.7 Best Combination

Here the best preprocessing combinations for the different feature extraction/matching methods regarding finger vein recognition are summarized. Figure 28 shows the effect of the best preprocessing combination found for SIFT on the two example finger vein images of figure 27. The original image is depicted on the left and the preprocessed one on the right. It can be clearly seen that the vein structure becomes more visible in both images.

- For SIFT / SURF: LeeRegion, HFE Filter, Circular Gabor Filter, Resize (0.5)
- For Maximum Curvature, Repeated Line Tracking, Huang Wide Line Detector and Template Matching: LeeRegion, Normalisation, Gabor Filter, Resize (0.5)
- For Local Binary Patterns: LeeRegion, Normalisation, Gabor Filter, Denoising, Resize (0.5)

It has to be noted that CLAHE was used in conjunction with the HFE Filter and the Circular Gabor Filter (as a replacement for the grey level grouping), so this is why CLAHE is not explicitly listed above.

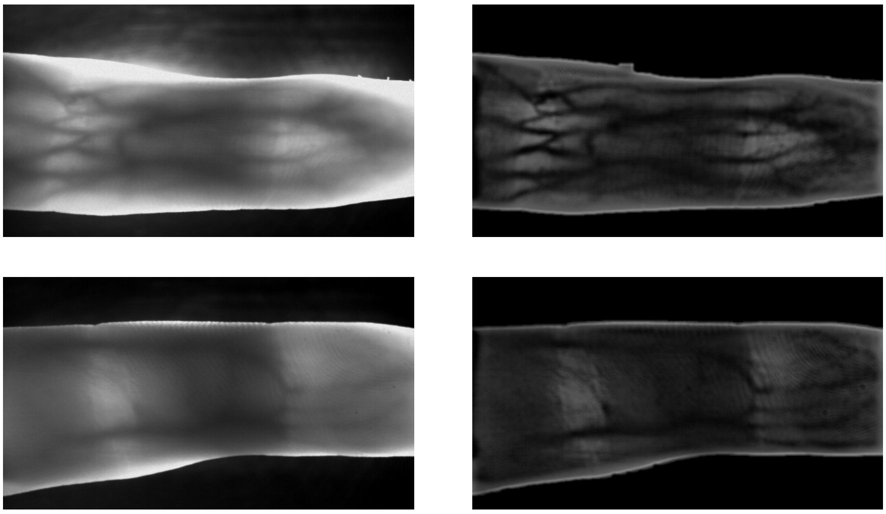

Fig. 28: Example of best preprocessing combination (for SIFT)

7.3 Feature Extraction and Matching Techniques

The first 3 techniques, which are discussed now and used during this thesis, basically aim to extract the vein pattern from the background using different approaches, which lead to a binary image. These binary images are then compared using a simple correlation measure. Also a custom implementation of a keypoint based technique using SIFT/SURF keypoints was used. Another feature extraction technique is an LBP based one, which uses the hamming distance for matching two feature images. None of these feature extraction techniques uses characteristics of the veins, e.g. vein crossings and endings, which are like minutiae in fingerprints. The last one is a simple template matching technique using adaptive local thresholding for binarisation of the input images and then again using a correlation measure for matching the images.

7.3.1 Maximum Curvature

This technique proposed by Miura et al. [37] emphasises the centre lines of the veins and is therefore insensitive to changes in the width of the veins. Processing of the finger vein image is done in three steps.

The first step is to extract the centre positions of the veins. F is the finger image, $F(x,y)$ is the intensity value of a pixel at position (x,y). $P_f(z)$ is the cross-sectional profile of $F(x,y)$ at any direction and position, where z is a position in the profile. To relate a position of $P_f(z)$ to $F(x,y)$ there is the mapping function T_{rs} defined as: $F(x,y) = T_{rs}(P_f(z))$. To find the centre positions, the local maximum curvature in the cross-sectional profile P_f in four directions, horizontal, vertical and the two oblique directions, based on the first and second derivatives is determined. The curvature $\kappa(z)$ can be represented as:

$$\kappa(z) = \frac{d^2P_f(z)/dz^2}{\{1 + (dP_f(z)/dz)^2\}^{\frac{3}{2}}}$$

Each profile is then classified as being concave or convex (curvature positive or negative). Only local maxima of $\kappa(z)$ in concave profiles are calculated and indicate the centre positions of the veins. The positions of these points are defined as z'_i where $i = 0, 1, ..., N-1$ and N is the number of local

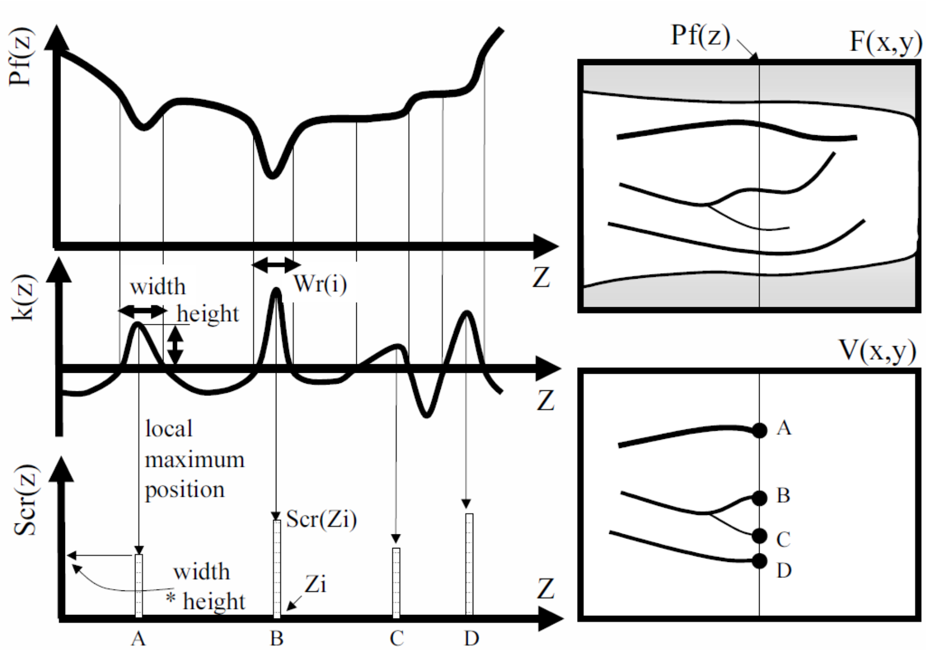

Fig. 29: Profile, curvature and probability score of veins [37]

maximum points in the profile. Each centre position is then assigned a score $S_{cr}(z)$, according to the width $W_r(i)$ and curvature $\kappa(z_i')$ of the region where z_i' is located.

$$S_{cr}(z_i') = \kappa(z_i') \times W_r(i)$$

The scores are assigned to a place V, which is the result of emphasising the veins:

$$V(x_i', y_i') = V(x_i', y_i') + S_{cr}(z_i')$$

where (x_i', y_i') represents a point defined by $F(x_i', y_i') = T_{rs}(P_f(z_i'))$. The relation of the cross-sectional profile, the curvature and the score can be seen in figure 29.

Afterwards the centre positions of the veins are connected. Due to noise or other distortions, some pixels may not have been classified correctly at the first step, so a filtering operation is applied to all pixels in all four directions.

$$C_{d1}(x, y) = min\{max\{V(x+1, y), V(x+2, y)) + max(V(x-1, y), V(x-2, y))\}\}$$

This filter just compares the central pixel with its two neighbours. If the central pixel has a small value but the neighbours have large values, it is assumed to be a vein and its value is increased to connect the line. If only the central pixel but none of its neighbours has a high value it is assumed to be noise and its value is decreased. If all pixels have high values nothing is changed. Subsequently the maximum of $C_{d1}, C_{d2}, C_{d3}, C_{d4}$ is selected at each pixel position and forms the so called locus space image $G(x, y)$.

Finally a binary image is generated from $G(x, y)$, representing the vein pattern using the median of the locus space as a threshold.

7.3.2 Repeated Line Tracking

This approach also proposed by Miura et al. [38] is based on tracking dark lines starting repeatedly at various random positions in the image. Veins appear as valleys in the cross-sectional profile of the image. At first the locus space T_r is initialized, which is a matrix with the same size of the input image, where every entry stores the number of times the corresponding pixel has been tracked as dark line. At each round the tracking point is initialized at a random position (x_c, y_c) and a locus position table T_c is initialized, where all tracking points found in the current round are stored. The tracking point is then moved pixel by pixel along a dark line, where the depth of the valley indicates the movement direction (the pixel is moved to where the valley is deepest, i.e. the value of $V_l = (s+t) - 2p$ has its maximum). Figure 30 shows the detection of a dark or vein line, respectively. The detailed steps are:

Defining the moving-direction attributes D_{lr} and D_{ud}. They are defined at the beginning of each round to bias the selection of the next tracking point towards a given direction, as veins tend to move in straight lines. As the tracking point is forced to go into that direction, this prevents the track from curving excessively.

$$D_{lr} = \begin{cases} (1,0) & if\, R_{nd}(2) < 1 \\ (-1,0) & else \end{cases}$$

$$D_{ud} = \begin{cases} (0,1) & if\, R_{nd}(2) < 1 \\ (0,-1) & else \end{cases}$$

where $R_{nd}(n)$ is a uniform random value between 0 and n.

The next step is to find a candidate new tracking point, if possible. Candidates are the neighbouring pixels, which have not been previously assigned as tracking point in the current round. Formally this can be written as:

$$N_c = \bar{T}_c \cap R_f \cap N_r(x_c, y_c)$$

with N_c as the set of candidate pixels for the new tracking point. $N_r(x_c, y_c)$ is the neighbouring pixels function, defined as:

$$N_r(x_c, y_c) = \begin{cases} N_3(D_{lr})(x_c, y_c) & if\, R_{nd}(100) < p_{lr} \\ N_3(D_{ud})(x_c, y_c) & if\, p_{lr} + 1 \leq R_{nd}(100) < p_{lr} + p_{ud} \\ N_8(x_c, y_c) & if\, p_{lr} + p_{ud} + 1 \leq R_{nd}(100) \end{cases}$$

with p_{lr} and p_{ud} as the probability that a horizontal or a vertical direction is chosen, respectively, N_8 is the eight-neighbourhood of the pixel (x_c, y_c) and N_3 is a set of 3 neighbouring pixels, which direction is determined by the moving direction attribute D:

$$N_3(D)(x,y) = \{(D_x + x, D_y + y), (D_x - D_y + x, D_y - D_x + y), (D_x + D_y + x, D_y + D_x + y)\}$$

where D can be either D_{lr} or D_{ud}.

The final step is to find the direction that has the deepest valley in the cross sectional profile. Therefore a line evaluation function is used:

$$\begin{aligned} V_l \quad = \quad & \max_{(x_i,y_i)\in N_c} F(x_c + rcos\theta_i - \frac{W}{2} sin\theta_i, y_c + rsin\theta_i + \frac{W}{2} cos\theta_i) \\ & + F(x_c + rcos\theta_i + \frac{W}{2} sin\theta_i, y_c + rsin\theta_i - \frac{W}{2} cos\theta_i) \\ & - 2F(x_c + rcos\theta_i, y_c + rsin\theta_i) \end{aligned}$$

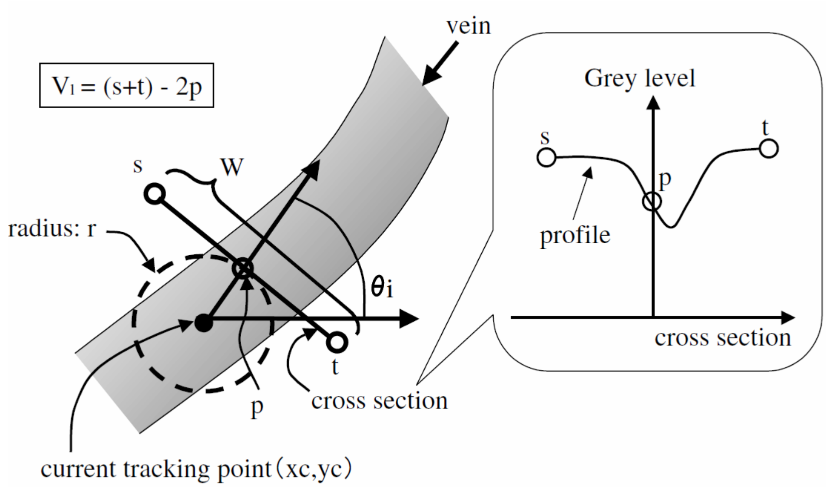

Fig. 30: Dark line detection [38]

where $F(x, y)$ is the pixel value at position (x, y) in the image, r is the distance between (x_c, y_c) and the cross section, θ_i is the angle between line segments $(x_c, y_c) - (x_c + 1, y_c)$ and $(x_c, y_c) - (x_i, y_i)$ and W is the width of the profiles in the cross section.

To speed up the process, only a certain number of angles θ_i is checked. If V_l is positive, the pixel is the new tracking point, it is added to T_c and used as starting point for the next iteration. If V_l is negative or zero, none of the neighbouring pixels are valid candidate points, i.e. they are not inside a vein or in other words there is no valley detected. In this case the current round is finished, T_r is updated by incrementing all values that are set to 1 in T_c and a new tracking operation is started at a random point.

The number of times each pixel is tracked as a dark line is recorded in the locus space matrix T_r. Pixels that are tracked multiple times as being a line statistically have a high likelihood of belonging to a blood vessel, i.e. they have a high value in locus space image. The tracking procedure is repeated N times. The higher N, the more accurate are the vein pattern results, but the computation time increases. The authors suggest to use at least a value of $N > 3000$ for good results. Finally again binarisation is applied to the locus space image the get the binary output vein image.

The parameters of this approach are N (number of starting times), W (valley width) and r (search radius).

7.3.3 Wide Line Detector

Huang et al. [2] suggested a quite simple method to extract the vein lines. This technique is essentially a local adaptive thresholding one using isotropic non-linear filtering, i.e. thresholding inside a neighbourhood region. Each pixel has a circular neighbourhood

$$N(x_0, y_0) = \{(x, y) | \sqrt{(x - x_0)^2 + (y - y_0)^2} \leq r\}$$

with a radius r. A centre pixel with its circular neighbourhood is depicted in figure 31. The difference of each pixel inside the neighbourhood to the central pixel is determined:

$$s(x, y, x_0, y_0, t) = \begin{cases} 0 & F(x, y) - F(x_0, y_0) > t \\ 1 & else \end{cases}$$

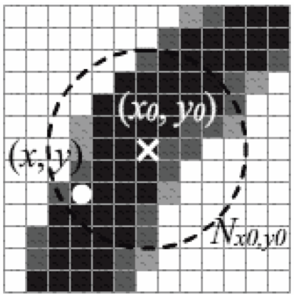

Fig. 31: Circular neighbourhood region [2]

Then the number of pixels inside this neighbourhood, which have a difference smaller than a set threshold, is determined:

$$m(x_0, y_0) = \sum_{(x,y)\epsilon N(x_0,y_0)} s(x, y, x_0, y_0, t)$$

This number is again thresholded to get a binary vein image:

$$V(x_0, y_0) = \begin{cases} 0 & m(x_0, y_0) > g \\ 255 & else \end{cases}$$

7.3.4 Matching using Correlation (Miura Matcher)

For matching the binary output images the approach used in [38] and [37] was adopted, which is essentially the calculation of the correlation between the input and reference image. As the input images are not registered to each other and only coarsely aligned (using LeeRegion), the correlation between the input image $I(x, y)$ and the reference one is calculated several times, while shifting the reference image $R(x, y)$ in x- and y-direction.

$$N_m(s, t) = \sum_{y=0}^{h-2c_h-1} \sum_{x=0}^{w-2c_w-1} I(s + x, t + y)R(c_w + x, c_h + y)$$

where $N_m(s, t)$ is the correlation value.

The maximum value of the correlation is normalised and used as matching score:

$$score = \frac{N_{m_{max}}}{\sum_{y=t_0}^{t_0+h-2c_h-1} \sum_{x=s_0}^{s_0+w-2c_w-1} I(x, y) + \sum_{y=c_h}^{h-2c_h-1} \sum_{x=c_w}^{w-2c_w-1} R(x, y)}$$

where s_0 and t_o are the indices of $N_{m_{max}}$ in the correlation matrix $N_m(s, t)$. The resulting score values are in the range of $0 \leq score \leq 0.5$.

7.3.5 Local Binary Patterns (LBP)

This feature extraction scheme is based on the use of local binary patterns (LBP). LBP compares the grey level of a centre pixel to its neighbouring pixels and assigns a value depending on the corresponding binary code resulting from the binary pattern of the neighbourhood pixels. The original LBP is a

3x3 non-parametric operator. It can also be defined as an ordered set of binary values determined by comparing the grey values of a centre pixel to its 8 neighbouring pixels:

$$LBP(x_c, y_c) = \sum_{n=0}^{7} s(i_n - i_c) \cdot 2^n, \quad s(x) = \begin{cases} 1 & if\, x \geq 0 \\ 0 & if\, x < 0 \end{cases}$$

I used a multi-scale version of LBP. Matching is done based on Hamming distance. The Hamming distance is basically calculated by applying XOR to the two feature images combined with an additional mask:

$$HD = \frac{||(codeA \otimes codeB) \cup maskA \cup maskB||}{||maskA \cup maskB||}$$

where $codeA$ and $codeB$ are the extracted local binary patterns, $maskA$ and $maskB$ are binary masks, where 1 means the pixel is inside a vein region and 0 means it is not. The masks are calculated using a 3×3 window in which the standard deviation of the pixels from the centre pixel is determined. If it is above a threshold (2), it is denoted as a vein region, otherwise not. Note that the extracted patterns are not binary, thus $\otimes$ is an XOR operator between the corresponding pairs of bits in the extracted patterns (binary representation of the integer values). $\cup$ is an OR operator. The Hamming distance values are in a range of 0...1 where 0 means a perfect match. Note that this matcher also performs a shifting of the codes (images), done to compensate misalignments of the input images, which is not stated in the equations above. Shifting is done in the same way as by the correlation matcher described above.

7.3.6 Template Matching

A basic adaptive binarisation approach [39] was used to assess the advanced binarisation techniques, which try to model the vein shape in their binarisation strategy. This technique uses adaptive thresholding for generating the binary vein image. The resulting binary images can then be compared against each other using XOR, Hamming distance or a correlation measure.

7.3.7 SIFT / SURF

Keypoint based techniques try to use information from the most discriminative points as well as considering the neighbourhood and context information of these points by extracting keypoints and assigning a descriptor to each keypoint. In this thesis SIFT [40] (Scale Invariant Feature Transform) was used instead of SURF [41] (Speeded Up Robust Features) because its performance was superior.

Especially in the case of SURF many strong keypoints are along the finger boundaries. As the classification should not be based on the finger shape but on the vein pattern and descriptors of these keypoints may also contain irrelevant background information, a filtering step was implemented at which all keypoints inside a window with a predefined width along the finger boundaries are discarded. Using the keypoints along the finger boundaries could lead to false or ambiguous matches. In addition a minimum number of keypoints can be defined and if after filtering only too few keypoints remain, the feature extraction is re-run with adapted parameters to extract more keypoints until at least the minimum number of keypoints are found. Note that there is no upper bound, i.e. no maximum number of keypoints defined.

SIFT and SURF matching is done based on the keypoint descriptors. The keypoint with the smallest distance to the reference keypoint is the matched one (nearest neighbour search) if the distance is below a certain threshold or otherwise the keypoint cannot be matched. A keypoint could have small distances to more than one other keypoint and this would lead to ambiguous matches, thus a ratio threshold is used: a match is only valid if the distance of the best matched keypoint is at least k times smaller than to the second best one (or all other points).

Matching returns a set of matched keypoints with associated distances. The simplest way to compute a final matching score is to use the number of matched keypoints only. A slightly better way is to use the ratio of matched keypoints to the maximum number of possible matches. There are also techniques involving the distances between the matched keypoints but they performed worse than the two simple techniques, so they have not been considered.

8 Experimental Setup

In this chapter details regarding the experimental setup to evaluate the impact of sensor ageing related pixel defects in fingerprint, finger- and hand vein images on the matching performance of different feature extraction and matching approaches are explained. At first the databases used for fingerprint, finger vein and hand vein evaluation, respectively, and also the test protocol are described. This is followed by the determination of the simulation parameters.

8.1 Fingerprint Database

8.1.1 Casia 2009 and 2013

In order to estimate the defect growth rate and the pixel defect parameters (defect offset and distribution of defect amplitudes) the images of the Casia Fingerprint Database[7] should have been used. There is one set of images which was acquired in 2009 and another set of images acquired in 2013, so there is a time span of 4 years in between, which should be enough for at least some sensor ageing related pixel defects to develop.

The 2009 data was captured using an URU4000B fingerprint sensor from Digital Persona[8]. There are a total of 3920 images acquired from 49 subjects in 3 sessions. The 2013 data was captured using three different fingerprint sensors, an URU4000B from Digital Persona, an URU4500 also from Digital Persona, which is the successor of the URU4000 and an Atmel Fingerchip T2. The images captured with the URU4000B and also the URU4500 are split in two directories, which indicates that there were two different sensors of this model used as I found out later. For the Atmel Fingerchip T2 there is only one directory but as this is a thermal sensor these images were not used during the experiments. There are a total of 1960 images acquired from 49 subjects for the URU4000B fingerprint sensor and also for the URU4500.

It later turned out that the images in 2009 and in 2013 were not captured with the same sensor. This means that an estimation of the defect growth rate and the defect parameters is not possible as the most important requirement, that the same sensor is used for both sets of images, is violated. Following is a proof by contradiction that the images have truly been captured with at least two different sensors.

8.1.2 Sensor Identification

It was not clear if the same sensor which was used in 2009 was again used in 2013, but in order to be able to determine the sensor's defect growth rate for the Casia data sets it is inevitable, that both image sets were taken using the same sensor. To check whether the sensor used in 2013 is the same as the one used in 2009 and vice versa, the sensor identification methodology of Höller and Uhl [42], which is based on the PRNU noise residuals of the sensor, was adopted.

The check was done using proof by contradiction so the assumption is that the two data sets were acquired using different sensors and it was tried to derive a contradiction to the assumption from the results.

The methodology works as follows: At first the PRNU noise residual of every image is extracted using 4 patches with a size of 128x128 pixels, located in the corners of each image. Then 850 images are chosen randomly from each data set which results in a total of 1700 images. 50 images of each of these 850 images per data set are used to calculate a PRNU fingerprint $\hat{K}$ for the data set and the remaining 800 images are used to calculate the normalised cross correlation (NCC) scores between the

[7] CASIA-FingerprintV5, `http://biometrics.idealtest.org/`

[8] They are now part of Crossmatch Technologies. No information regarding the URU4000B fingerprint sensor could be found on their new website `http://www.crossmatch.com/`.

	1	**2**	**3**	**4**	**5**	**6**	**7**
finger_uru4000 (**1**)							
finger_uru4000_finger0358 (**2**)	49.9722						
finger_uru4000_finger1267 (**3**)	49.8889	49.3056					
uru4000_1 (**4**)	0	0	0				
uru4000_2 (**5**)	0	0.0556	0	0			
uru4500_1 (**6**)	0	0.0556	0	0	0		
uru4500_2 (**7**)	0	0	0	0	0	0	

Tab. 3: Cross-sensor matching results

fingerprints and the noise residuals. As there are two data sets with 800 remaining images each, this leads to 800 matching and 800 non-matching scores, at which the fingerprint $\hat{K}$ comes from the same sensor as the images and the fingerprint $\hat{K}$ comes from a different sensor as the images, respectively. Afterwards, the equal error rate (EER) is calculated using all correlation values ρ by comparing the two data sets. A confidence interval (CI) of 95% is used to estimate the real variability of these values. For this purpose the EER and the corresponding threshold are calculated 1000 times on a set of m matching and n non-matching ρ values. The m matching correlation values and the n non-matching correlation values are drawn from the matching and non-matching data set, respectively, using sampling with replacement. The range which contains 95% of the total 1000 values (i.e. 950 values then) is the interval of confidence.

If the same sensor has been used for both data sets, the resulting EER would be very high, i.e. nearly 50%, which means the matching and non-matching NCC scores are almost identical. So the matching scores cannot be distinguished from the assumed non-matching ones, because the PRNU fingerprints $\hat{K}$ are present in the images from the same as well as in the images from the other data set. This would be in contradiction to the assumption that different sensors have been used and therefore the same sensor must have been used for both data sets.

Results for Casia 2009 and 2013 Database The results for the cross-sensor matching are shown in table 3. The Casia Database consists of 2 directories, one for 2009 and one for 2013. The first one contains 3 sub-directories: finger_uru4000 (**1**), finger_uru4000_finger0358 (**2**) and finger_uru4000_finger1267 (**3**). The second one contains 5 sub-directories: Finger_T2, uru4000_1 (**4**), uru4000_2 (**5**), uru4500_1 (**6**) and uru4500_2 (**7**). The directory named Finger_T2 was not evaluated because the images were captured with a thermal sensor and not with one of the URU sensors. Cross matching was done between all the different directories. As it can be seen from the results, the EER is high, i.e. nearly 50%. Thus the images in the directories finger_uru4000, finger_uru4000_finger0358 and finger_uru4000_finger1267 were captured using the same sensor. As all other determined EERs are 0 or close to 0, thus none of the images in the other directories have been captured with the same sensor, especially the images from 2013 have all been captured with different sensors than the ones in 2009. Thus it does not make sense to try to detect sensor ageing related pixel defects inside the images as they do not have a temporal relationship. There were no other data sets of the sensors used with a timelapse in between available to be able to estimate the approximate defect growth rates and parameters directly from the images. So I used an empirical formula to estimate the defect growth rate based on the image sensor characteristics as described in section 8.4.1.

	Sensor Type	Model	Image Size	Resolution
DB1	Optical	CrossMatch V300	640×480	500 dpi
DB2	Optical	Digital Persona U.are.U 4000	328×364	500 dpi
DB3	Thermal Sweep	Atmel FingerChip	300×480	512 dpi
DB4	Synthetic	SFinGe v3.0	288×384	500 dpi

Tab. 4: Databases of the FVC2004 data set

8.1.3 FVC2004 Dataset

For the ageing simulation experiments in this work, the example fingerprint images from the FVC2004 [43] data set were used and also the test procedure was adopted. The FVC (Fingerprint Verification Competition) was a contest organised by the Biometric System Laboratory of the University of Bologna, the Biometric Test Centre of San Jose State University and the Biometrics Research Lab - ATVS of the Universidad Autonoma de Madrid comparing the performance of different fingerprint matching algorithms. Further information on the FVC can be found on their website[9].

The data set consists of three databases containing fingerprint images, acquired with different sensors and one containing synthetic fingerprint images. The images are provided in uncompressed TIFF format.

Each database consists of two sets of images:

- Set A consists of 100 fingers with 8 images per finger, which is a total of 800 images. This one is only used in the final experiments.
- Set B consists of 10 fingers with 8 images per finger, which is a total of 80 images. This set is used for preliminary tests and parameter tuning.

As the fourth database (DB4) contains synthetic images and in the present thesis the impact of sensor ageing related defects should be investigated, DB4 was not used because the images are not acquired using an image sensor. DB3 was acquired using a thermal sensor and it is not clear how this type of sensor is affected by ageing, thus DB3 was also not used. So only DB1 and DB2, which were both acquired using an optical sensor, were used during the experiments.

The implementation of the non-minutiae based matchers has a rather low processing speed, thus the images of DB1 were resized to 512×512 pixels because a 2D Fast Fourier Transform is used, which can only be applied to images that have dimensions of a power of 2, i.e. the images of DB1, having an original resolution of 640×480, would have been automatically resized to 1024×512 pixels, which would further extend the processing time. Resizing was done by first cropping the images on the left and right side, i.e. a stripe of 64 pixels width was removed on the left and right side leaving the centre area and resulting in an image size of 512×480. Afterwards the top and bottom of the image was extended by a stripe which is 16 pixels high and its pixels were set to white (255), which corresponds to the background of the original images.

Test Procedure The test procedure of all the FVC contests basically consists of two types of tests:

- Genuine Tests: To determine the False Non Match Rate (FNMR) the genuine tests have to be carried out. Each sample image is matched against all remaining samples of the same finger. No symmetric matches are performed. Based on the Set A of databases from the FVC2004 data

[9] `http://bias.csr.unibo.it/fvc2004/`

Fig. 32: Example images of the FVC2004 database (left: DB1, right: DB2)

set, containing 100 fingers with 8 images each, the number of genuine tests is:

$$\frac{n_{images} \cdot (n_{images} - 1)}{2} \cdot n_{fingers} = \frac{8 \cdot 7}{2} \cdot 100 = 2800$$

- Impostor Tests: These tests are used to determine the False Match Rate (FMR). Therefore, the first imprint of each finger is matched against the first imprint of all the remaining fingers. As with the genuine tests, no symmetric matches are performed. Again based on the Set A of the databases, the number of impostor tests is:

$$\frac{n_{fingers} \cdot (n_{fingers} - 1)}{2} = \frac{100 \cdot 99}{2} = 4950$$

8.2 Finger Vein Database

For the evaluation I used the University of Twente Finger Vascular Pattern Database (UTFVP) [33], kindly provided by R.N.J Veldhuis. This data set consists of a total of 1440 images, taken from 60 persons, 6 fingers per person (index, ring and middle finger of each hand) and 4 images of each finger. The images were acquired in 2 sessions with a time lapse of 15 days between the sessions. Each finger was captured twice during one session. 73% of the subjects were male and 82% were right handed. The images have a resolution of 672 x 380 pixels, a density of 126 pixels/cm and are stored in 8 bit grey scale PNG format. The width of the visible blood vessels is 4 - 20 pixels.

8.2.1 Test Procedure

The original test procedure from [33] was finally not used during this work due to its high number of impostor matches. To determine the EER at first the positive and negative match scores have to be determined, which is similar to the determination of the FNMR and FMR, respectively. For computing the FNMR and FMR the test protocol of the FVC2004 was adopted:

- Genuine Matches: Each image of each finger is compared with all remaining images of the same finger. No symmetric matches are performed. The total number of positive compares is therefore:

$$\frac{4 \cdot 3}{2} \cdot 360 = 2160$$

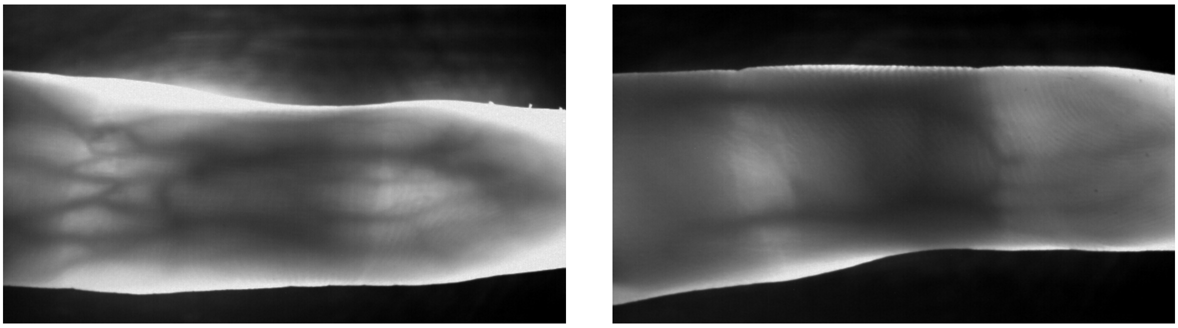

Fig. 33: Example images of the UTFVP finger vein database

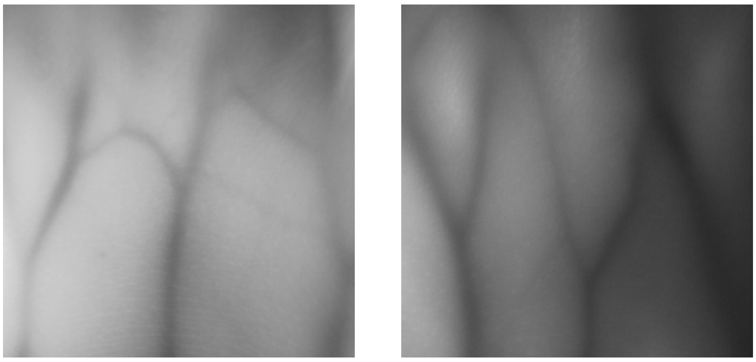

Fig. 34: Example images of the hand vein database

- Impostor Matches: Using the same test procedure as in the FVC2004, i.e. the first image of each finger is compared against the corresponding first image of the same finger of all remaining subjects. This results in a total of:

$$\frac{60 \cdot 59}{2} \cdot 6 = 10620$$

8.3 Hand Vein Database

The hand vein image database was collected at the University of Salzburg using a custom made prototype hand vein scanner. From the raw images taken by the scanner a ROI (region of interest) was extracted, which is roughly located in the centre of the hand surface and has an area of 500×500 pixels. There are a total of 107 subjects with at least 1 image per hand per subject. For most of the subjects only images of the left hand were captured, but for some subjects images of the right hand were captured too. The images are split in reflected light and transillumination images, resulting in a total of 620 images for transillumination and 593 for reflected light. The images are stored in JPEG format with a resolution of 500×500 pixels.

Custom Subset As stated above, there is not the same number of images available for each subject (hand) in the database. There can be anything from 1 to 9 images per hand. To simplify the test protocol, a custom subset out of the transillumination images contained in the database was used. This subset consists of 100 subjects (hands) with exactly 3 images per subject, thus a total of 300 images. Only transillumination images were used, because the achieved EER is higher compared to using reflected light images only and also to using both, transillumination and reflected light images.

8.3.1 Test Procedure

Again for the hand vein experiments the test protocol of the FVC2004 was adopted.

- Genuine Tests: Each sample image is matched against all remaining samples of the same hand. No symmetric matches are performed. Based on the the custom subset used for evaluation, containing 100 hands with 3 images each, the number of genuine tests is:

$$\frac{n_{images} \cdot (n_{images} - 1)}{2} \cdot n_{hands} = \frac{3 \cdot 2}{2} \cdot 100 = 300$$

- Impostor Tests: Following the FVC2004 test protocol the first image of each hand is matched against the first image of all remaining hands. As with the genuine tests, no symmetric matches are performed. Again based on the custom subset, the number of impostor tests is:

$$\frac{n_{hands} \cdot (n_{hands} - 1)}{2} = \frac{100 \cdot 99}{2} = 4950$$

8.4 Simulation Settings

As I was not able to determine the number and parameters of the sensor ageing related pixel defects directly from images captured by the respective sensor due to the lack of usable data, i.e. I had no two data sets acquired by the same sensor with a timespan of several years in between, the defect growth rate could not be estimated directly. Thus, I used the formula of Chapman et al. [9] to estimate the defect growth rate for the different sensor types according to the parameters of the particular image sensor.

8.4.1 Defect Growth Rate

Fingerprint Scanner As I showed in 8.1.2 the Casia data sets of 2013 and 2009 were captured using different sensors and thus it was not able to derive the defect growth rate and defect parameters from these data sets. So I had to estimate the defect growth rate using the following formula from Chapman et al. [9], explained in 3.7 :

$$D = A \cdot S^B \cdot ISO^C$$

As I have not got technical information from Digital Persona[10] regarding their U.are.U4000B and U.are.U4500 fingerprint sensors, I did some research on commonly used image sensors for that purpose. The most widely used ones are CMOS sensors with a pixel sizes between $5.0\mu m$ and $7.4\mu m$ (square pixels).

I also ordered an U.are.U4000B fingerprint sensor and disassembled it to have a look at the image sensor. There was no label on it but I was able to measure the sensor size. In figure 35 a picture of the image sensor can be seen. The outside dimensions (dark yellow square) are $10.6 \times 10.6\,mm$ or about $410 \times 410\,pixels$ in the image. The actual image sensor area (red to yellow rectangle in the left and blue rectangle in the right picture) is about $77 \times 66\,pixels$. This means that the sensor's dimensions are $1.99 \times 1.71\,mm$ which corresponds to a sensor area of $3.404\,mm^2$.

The images captured by the U.are.U 4000b and U.are.U 4500 have a resolution of 356×328 pixels, thus, with the measured sensor size of $3.404\,mm^2$ the pixel size is: $\frac{1.99}{356} \times \frac{1.71}{328} = 5.59 \times 5.21\,\mu m$. I

[10] `http://www.crossmatch.com/UareU4500Reader/`

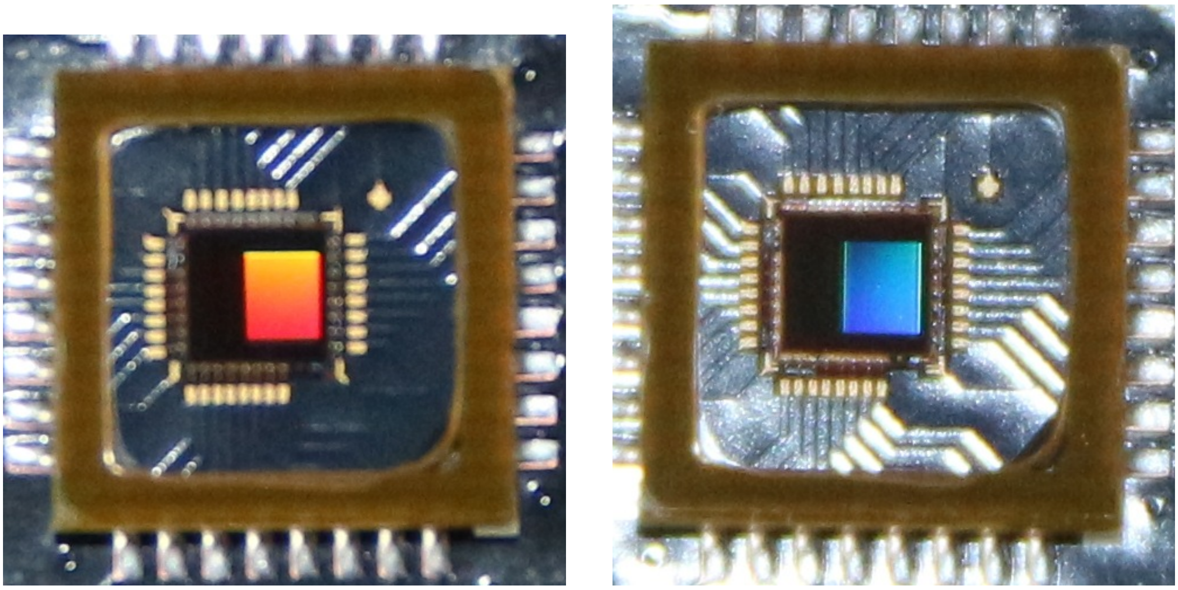

Fig. 35: Image sensor inside the U.are.U4000B fingerprint scanner

therefore assumed a pixel size of $5.4\,\mu m$. I further assumed that the images are taken at an ISO Level of 400. This results in a defect growth rate of:

$$D = 0.0742 \cdot 5.4^{-3.07} \cdot 400^{0.5} = 0.008375 \; defects/year/mm^2$$

or for the assumed sensor area of $3.404 mm^2$:

$$0.0285\, defects/year = 0.244 \; defects/MP/year$$

Finger Vein Scanner The images in the UTFVP database [33] were captured using the BCi5 monochrome CMOS camera produced by C-Cam Technologies. According to the data sheet of the manufacturer[11] this camera uses an CMOS image sensor with an active sensor area of $8.58 \times 6.86\, mm$, a pixel size of $6.7 \times 6.7\,\mu m$, a fill factor of 50% and a resolution of 1280×1024 pixels. There is no information available at which ISO level the images were captured. Therefore I assumed ISO level 400 again. This results in a defect growth rate of:

$$D = 0.0742 \cdot 6.7^{-3.07} \cdot 400^{0.5} = 0.00432 \; defects/year/mm^2$$

or for a sensor area of $58.859\, mm^2$:

$$0.254\, defects/year = 0.194 \; defects/MP/year$$

As the images contained in the database have a resolution of 672×380 pixels, the effective pixel defect rate is:

$$\frac{672 \times 380}{1280 \times 1024} \cdot 0.254 = 0.0495 \; defects/year$$

[11] www.c-cam.be/doc/Archive/BCi5.pdf

Hand Vein Scanner The custom build hand vein scanner uses a Canon EOS 5D Mark II DLSR camera as image capturing device. According to Canon this camera uses a CMOS image sensor with a size of $36.0 \times 24.0\,mm$ and a resolution of 5616×3744 pixels. Thus it has a pixel size of $6.41 \times 6.41\,\mu m$. The images were taken at ISO level 800. This results in a defect growth rate of:

$$D = 0.0742 \cdot 6.41^{-3.07} \cdot 800^{0.5} = 0.006997\ defects/year/mm^2$$

or for a sensor area of $864\,mm^2$:

$$6.045\,defects/year = 0.287\ defects/MP/year$$

According to the literature [6] DSLRs develop between 0.5 and 10 defects per year at ISO Level 400 so a value of 6.045 defects per year seems reasonable.

As only a ROI (region of interest) with a size of 500×500 pixels is used the effective pixel defect rate is:

$$\frac{500 \times 500}{5616 \times 3744} \cdot 6.045 = 0.0719\ defects/year$$

8.4.2 Hot and Stuck Pixel Amplitudes

Jessica Fridrich [18] and also others found out that there are stuck low, stuck high and in general stuck pixels with any value between 0 and 255 (for 8 Bit images). Thus, I assumed that stuck pixels can appear with any value between 0 and 255, uniformly distributed. According to Chapman et al. [6] but also Albert Theuwissen [3] the additional offset I_{Offset} of hot pixels or the dark current value, respectively follows an exponential distribution, i.e. hot pixels with a lower amplitude are more likely to occur.

Estimating the Exponential Distribution To be sure to use an accurate model for the distribution of the hot pixel amplitudes it has to be verified whether they are really exponentially distributed and then a fitting exponential distribution with the parameter μ has to be determined. First of all, the actual values of the hot pixel amplitudes and the number of hot pixels having these amplitudes have to be determined. As basis for that figure 3 (a) in [5], at an exposure time of $1/2\,s$ and for ISO Level 400, was used to get the reference data. Chapman uses an I_{offset} value, which is according to the pixel model in section 3.4.2, the dark response of a hot pixel, i.e. an additional offset to the electrical current induced by the incident light, depending on the dark current amplitude, the exposure time and the ISO level. For each database or set of images I tested, all images were captured using the same exposure time, ISO level and roughly the same environmental conditions like temperature, i.e. the dark current amplitude should remain constant. So instead of an I_{offset} value I used the pixel value offset directly for the simulation of pixel defects. As pixel values between 0 and 255 are used, an I_{offset} value of 0 simply corresponds to a pixel offset value of 0 and an I_{offset} value of 1 corresponds to a pixel value of 255.

As second source of reference data figure 6 from [3] was used. There the number of hot spots per sensor per day versus the hot spot amplitude in DN (digital numbers) is shown. I assumed that a DN value of 4500 corresponds to a pixel value of 255 and used the corresponding values as pixel offset values then. The values from the two figures were extracted as follows:

A suitable exponential distribution with its parameter μ was determined using MATLAB's built-in graphical distribution fit tool `dfittool`.

For the Chapman data which is given as a histogram, a bin width of 0.2 was used for the input data and a corresponding exponential distribution with $\mu = 0.1438$ at a standard error of 0.00045437 was found.

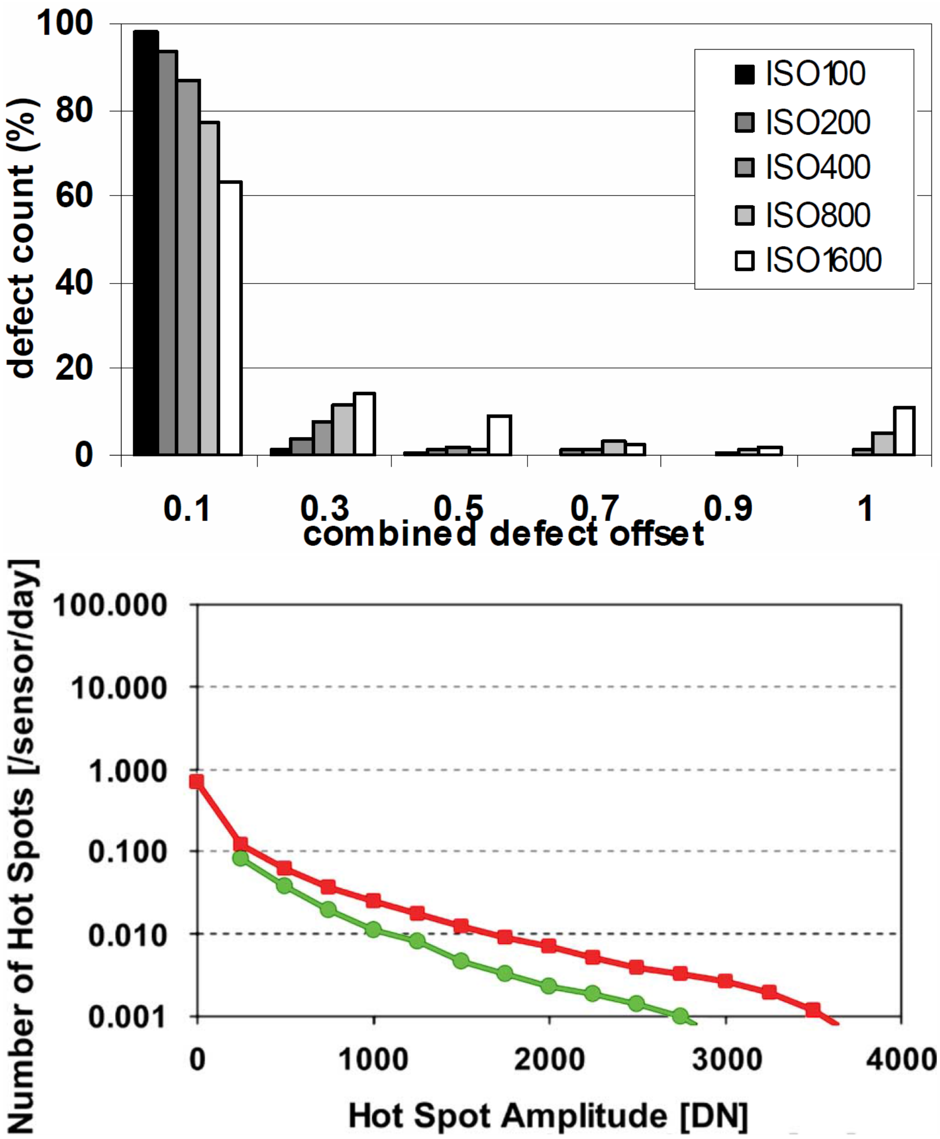

Fig. 36: Hot pixel amplitudes, top: Chapman et al. [5], bottom: Theuwissen, Albert [3]

I_{offset}	defect count (%)
0.1	87.3
0.3	8.3
0.5	1.9
0.7	1.2
0.9	0.2
1	1.1

Hot Spot Amplitude (DN)	Number of hot spots
0	0.6856
250	0.1208
500	0.06122
750	0.03475
1000	0.02569
1250	0.01762
1500	0.01254
1750	0.00893
2000	0.00712
2250	0.005264
2500	0.003892
2750	0.003346
3000	0.002668
3250	0.001973
3500	0.00119

Tab. 5: Defect counts of hot pixels with a certain amplitude, left: Chapman data, right: Theuwissen data

For the Theuwissen data at first the DN values were normalised by dividing each value by 4500 to get values in the interval of $[0, 1]$. But at first no suitable exponential distribution was found, so the first value which corresponds to $DN = 0$ and actually means that it is no hot pixel at all (having a hot spot amplitude of 0) was omitted. Still no exponential distribution could be found, so I interpreted the diagram as a histogram with a bin width of $\frac{250}{4500} = 0.05556$ and then best fit was obtained using $\mu = 0.16839$ with a standard error of 0.000961.

8.4.3 Simulation Parameters

The resulting defect rates according to the estimations are rather low, ranging from 0.028 to 0.072 defects per year, i.e. less than 1 defect per year. Statistically, even in 30 years there would only be 1 or 2 defective pixels visible in the finger or hand images. I started the simulations using a defect rate of 1 defect per year, with a simulation timespan of 30 years to get realistic results (no one would use a camera or biometric sensor for more than 30 years). The defect rate was doubled each time, i.e. 2 defects per year, 4 defects per year and so on, and the simulation was run again until there were noticeable effects on the recognition accuracy. As the defects were only barely noticeable below 1000 defects contained in a single image, I modified the simulation parameters and ran the simulation with a defect rate of 1000 defects per year for a time span of 10 years, which results in a total of 10000 hot or stuck pixel defects or 20000 hot and stuck pixels combined, respectively. Note that using both, stuck and hot pixel defects, the defect rate is twice as high as using a single defect type. The results only show the total number of defects, i.e. the combined defect rate is given.

The fingerprint images of the FVC2004 DB1 and DB2 have different resolutions. Also the finger vein and hand vein image resolutions are different from that of the fingerprint images. If a defect rate of 1000 defects per year is used, this results in a higher defect density for images with a smaller resolution and a lower defect density for images with a higher resolution, respectively. A higher defect density has a higher impact on the recognition accuracy. To be able to compare the results among each other, a defect rate per year per MP was used. The simulations were done for finger vein images first using a defect rate of 1000 defects per year. These images have a resolution of $672 \times 380 = 0.255\,MP$, thus for the other simulations a defect rate of $4000\, defects/MP/year$ was used.

The stuck pixel amplitudes were drawn uniformly out of the interval $[0, 255]$ and the hot pixel amplitudes were drawn according to an exponential distribution with $\mu = 0.15$.

All simulations were run using only stuck pixel defects, only hot pixel defects and then both, stuck and hot pixel defects, respectively.

During the tests two sets of images were used: the gallery database, which contains the unaged images and the probe database, which contains the aged images. In other words the template images were not aged, only the probe images were. This means that in each pair of matched images only one is aged, the other one is the original image from the data set. In addition for some of the experiments, also the template images were aged using the same defective pixels as the probe images to be able to examine the effects. In the case of gallery images aged, this is denoted with *TA* (templates aged) in the diagrams and in the text. An application scenario where also the gallery images are affected by sensor ageing is the use of an “old” biometric sensor, i.e. a sensor that was in use for at least a few years, for establishing a biometric database. In this case the probe images will contain ageing related pixel defects which will become more severe with increasing sensor age. In addition also the gallery images will contain some ageing related pixel defects, depending on the point in time when the database was established. For a thorough investigation of course all combinations of different progress in the ageing process should be tested. E.g. the database could have been established in year 1 and the probe images could be captured in year 12, or the database could be established in year 10 and the probe images also captured in year 10. Due to time restraints this would have been beyond the scope of this thesis so only the combinations where the gallery and probe images are captured at the same point in time, i.e. show the same ageing related pixel defects, were tested.

9 Results

This section presents the experimental results from running the ageing simulation for the different biometric traits, i.e. fingerprint, finger vein and hand vein, and also the different matchers used for each trait aimed to investigate the impact of sensor ageing related pixel defects on the recognition performance of each of the feature extraction and matching schemes used for the different biometric traits. More information on the experimental settings, especially the simulation parameters, the test data sets and the test protocol, can be found in section 8. The following subsection introduces some abbreviations for the evaluated feature extraction and matching schemes. After that each subsection deals with a specific biometric trait, beginning with finger veins, followed by hand veins and finally fingerprints. In each of these three subsections first of all the tested schemes are listed and specific settings or tests are described, if any. In addition some sample aged images for each trait are given. Afterwards the experimental results are provided in form of a table, stating the baseline EER and the EERs for the aged images, and diagrams showing the EER at different numbers of defective pixels for each of the evaluated approaches. This is followed by a discussion of the results. Not only the results of each individual approach are discussed but they are also compared to the other approaches that were tested. In addition I also try to give some explanations why some approaches are influenced whereas other are not and how to reduce the impact in general.

9.1 Abbreviations

The following list shows the abbreviations for the different feature extraction and matching schemes used in the diagrams and the text.

- MC: Maximum Curvature, described in section 7.3.1. Feature extraction method which was used for finger vein and also for hand vein evaluations.
- SIFT: SIFT based feature extraction method, also used for finger vein and hand vein evaluations, described in section 7.3.7.
- WLD: Huang Wide Line Detector, described in section 7.3.3. This is the second feature extraction method which is used for finger vein and hand vein experiments.
- RLT: Repeated Line Tracking, feature extraction method used only for finger vein experiments due to its high runtime and inferior results on hand vein images. Described in section 7.3.2.
- LBP: Local Binary Patterns, another feature extraction method which is described in section 7.3.5. This was only used for finger vein images due to its inferior results on hand vein images compared to MC and SIFT.
- AB: Adaptive Binarisation or Template Matching as the last feature extraction method which was used for finger vein and hand vein experiments. Details are described in section 7.3.6.
- NBIS: Means the mindtct minutiae extractor and bozorth3 matcher from the NIST Biometric Image Software package. Details can be found in section 6.4.3. NBIS was used for fingerprint images only.
- VF: VeriFinger matcher from Neurotechnology. This is a commercial package of a minutiae based feature extractor and matcher for fingerprint images. Details can be found in section 6.4.3. A trial version was used during the fingerprint evaluations.
- POC: Phase Only Correlation, a correlation based matcher for fingerprint images. A custom implementation was used, details are described in section 6.4.1.

- FC: Finger Code is a ridge feature based matcher for fingerprint images. Details of the custom implementation which was used during the evaluations can be found in section 6.4.2.

9.2 Finger Vein Results

In this subsection the results of the finger vein experiments are lined out. MC, RLT, WLD, LBP, AB and the SIFT based approach were tested. The simulations were run for hot pixels only, stuck pixels only and combined hot and stuck pixels. All simulations were run 5 times, i.e. 5 test sets with aged images are created. All the feature extraction and matching schemes are evaluated for each test set. The sensor defects are randomly distributed over the image and also the amplitudes or hot pixel offsets, respectively, are random values. Due to the random nature of the defect locations and values, some defective pixels may have a larger influence on the recognition accuracy than others, e.g. if there are several hot pixels inside a vein. Running the simulations 5 times should mitigate these effects but, as it can be seen on the results, not completely as there are still some statistical fluctuations. The mean of the EER is used as final result.

In addition, all evaluations were done using a denoising filter (median filter followed by an adaptive Wiener filter) during the preprocessing step, which reduces the sensor ageing related impact to a great amount even for those schemes which are heavily influenced. One drawback is that the baseline EER rises, i.e. the recognition accuracy drops if denoising is used. Also note that the results for LBP are the same with and without denoising as LBP uses denoising by default.

9.2.1 Sample Aged Images and Corresponding Feature Extraction

In figure 37 some sample finger vein images with no defects, 1000 and 10000 pixel defects and the corresponding region of the feature extraction image for MC, RLT and WLD are shown. One can clearly see that for MC due to the pixel defects the lines get broken and an additional line appears at 10000 defects which is clearly not a vein. The vein lines do not get broken in the case of RLT but they appear wider with some additional noise at the vein boundaries. WLD does not show much useful information any more at 10000 defects due to the noise caused by the defective pixels which appears as small circles inside the image.

9.2.2 Simulation Results

Table 6 shows the baseline EER for all the tested approaches with and without denoising. As stated above, in the case of MC, RLT and SIFT the baseline EER rises when denoising is used. In addition also the EER for 10000 defects present in the images (5000 hot and 5000 stuck pixels) are shown with and without denoising. Furthermore the percental rise of the EER compared to the corresponding baseline EER is given right next to each absolute EER value. By comparing the baseline EER values of the different approaches, it can be clearly seen that MC performs best with an EER of 0.006. The second best performing approach is RLT closely followed by SIFT, both achieving an EER of 0.02 which is more than 3 times higher than MC. On the fourth place is WLD with an EER of 0.031 followed by AB with 0.036. The worst performing approach is the LBP based one with an EER of 0.063, which is more than 10 times higher than the best performing approach, MC.

Figure 38 shows that not only SIFT and MC but also the simple AB scheme (presumably because it additionally relies on the finger outline and not solely on the vein structure) is hardly influenced at all by hot pixel defects, even if there are 10000 defects inside the image. The EER of MC rises to 0.0074 at 10000 defects but it is still the best performing approach. The EER of RLT doubles at 10000 defects but stays still below 0.045 so it is on the fourth place while SIFT gets on the second one with an EER of 0.024. The EER of AB rises to 0.043, which is a bit better than RLT and on the third place regarding the matching performance. WLD, as it is quite a simple thresholding method, is

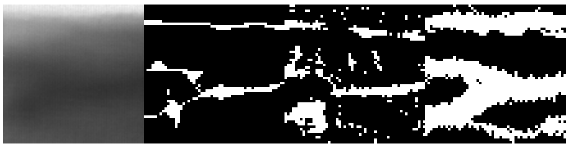

(a) No defects (b) MC (c) WLD (d) RLT

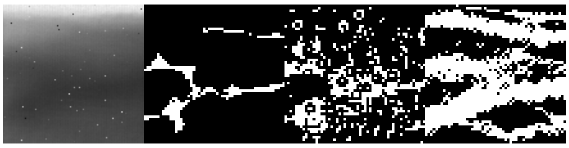

(e) 1000 defects (f) MC (g) WLD (h) RLT

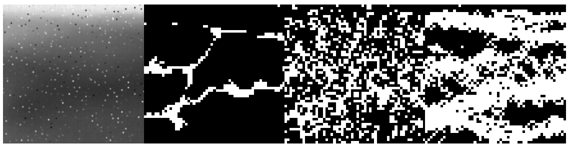

(i) 10000 defects (j) MC (k) WLD (l) RLT

Fig. 37: Sample aged images and feature extraction

Method	EER	EER	EER (%)	10000 Defects	EER (%)	10000 Defects	EER (%)
	Baseline	w. denoising				w. denoising	
MC	0.006	0.010	60.5%	0.017	184.4%	0.010	5.7%
RLT	0.020	0.021	4.2%	0.077	285%	0.021	-0.4%
WLD	0.031	0.025	-16.8%	0.361	1079%	0.028	11.9%
LBP	0.063	0.063	0%	0.068	8.1%	0.068	8.3%
SIFT	0.020	0.022	11.6%	0.029	46.1%	0.023	1.7%
AB	0.036	0.036	-1.5%	0.048	30.6%	0.036	-1.1%

Tab. 6: Finger-vein EER baseline and at 10000 defects

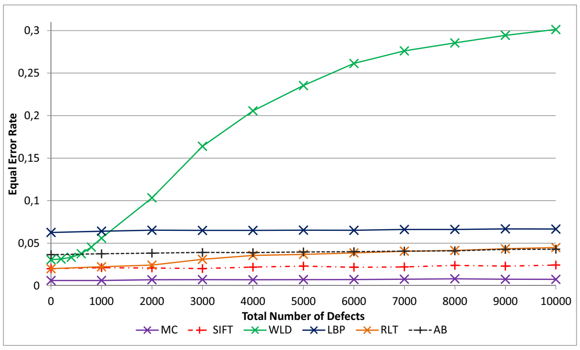

Fig. 38: EER hot pixels only

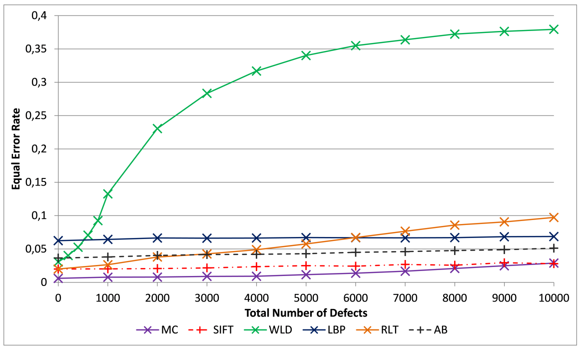

Fig. 39: EER stuck pixels only

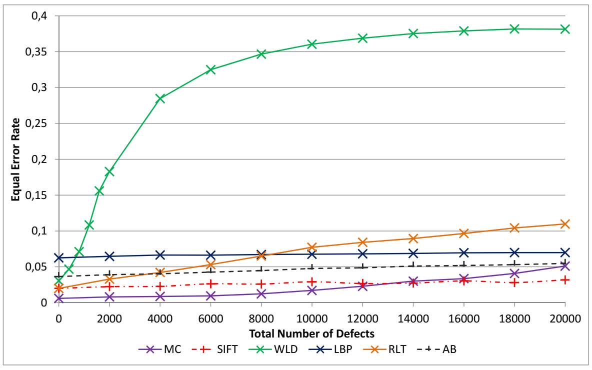

Fig. 40: EER hot and stuck pixels

influenced dramatically. At 10000 defects the EER is 0.301 and thus ten times higher than its baseline EER. Hence it is the worst performing approach. The performance of LBP is hardly influenced at all by the noise introduced due to sensor ageing, even in the case of 10000 defects, its EER is 0.067 which corresponds to an increase of 6.5%. LBP uses denoising by default but it is still no better than the fifth place regarding the performance ranking.

In figure 39 it can be seen that the influence of stuck pixels is more severe than of hot pixels. The average amplitude or hot pixel offset, respectively, is quite low thus the original grey value is only slightly changed. Stuck pixels always have the same fixed value, independent of the original grey value, so they are actually more likely to “break” the vein lines inside the image.

SIFT is again quite stable with an EER of 0.028 at 10000 hot pixel defects, which is now placed first in the performance ranking. MC is affected more by stuck than by hot pixels and at 10000 defects the performance of MC gets slightly worse than the SIFT one with an EER of 0.029, which means it is the second best approach. Also the performance of AB is influenced more by stuck pixels than by hot ones as the EER rises to 0.051 at 10000 stuck defects in comparison to 0.043 for hot ones but it is again on the third place. RLT is influenced more than MC and AB as its EER rises to 0.097 in comparison 0.045 at 10000 hot pixel defects and thus it is on the fourth place. In accordance with MC, AB, RLT and SIFT, WLD is also influenced more by the stuck pixel defects. The EER of WLD rises to 0.379 at 10000 stuck pixel defects, which is again the last place regarding matching performance. Only for LBP the impact of stuck pixels is not significantly higher than for hot pixels due to denoising. Its EER for 10000 stuck pixels is 0.069 and it is on place four in the ranking this time.

If both, hot and stuck pixel defects, are present in the images, similar results as with the single defect types can be seen. WLD is influenced most and again the worst performing approach, whereas the performance of SIFT is quite stable, even if 20000 defects in total (10000 hot and 10000 stuck pixels) are present, i.e. the EER rises to 0.032. SIFT is the best performing approach at 20000 defects. The performance of MC gets worse than SIFT at 14000 defects (MC: 0.03 and SIFT: 0.027) and at 20000 defects it is the second best performing approach, only slightly ahead of AB which is on the third place. RLT is affected more than by hot pixels only but not as much as by stuck pixel defects only. At 20000 defects it performs worse than LBP and is the second worst performing approach. Its EER for a total of 10000 defects is 0.077 and at 20000 defects it is 0.111. The influence on LBP and AB is again negligible. The EER of LBP at 20000 defects is 0.07, thus it is on the fourth place in the performance ranking. AB achieves an EER of 0.054 at 20000 defects, which means it is the third best performing approach.

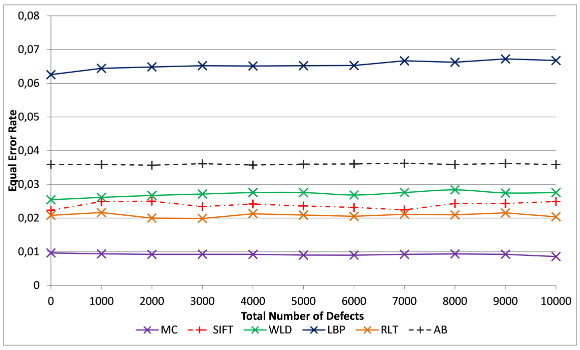

Fig. 41: EER hot pixels only with denoising

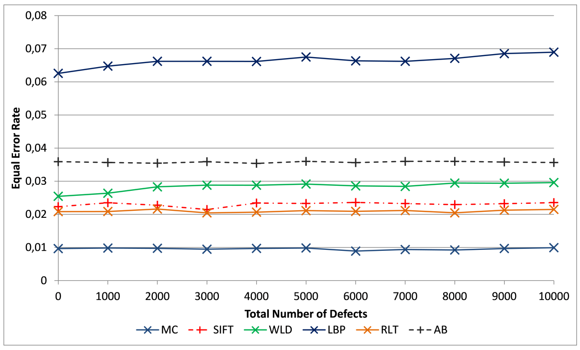

Fig. 42: EER stuck pixels only with denoising

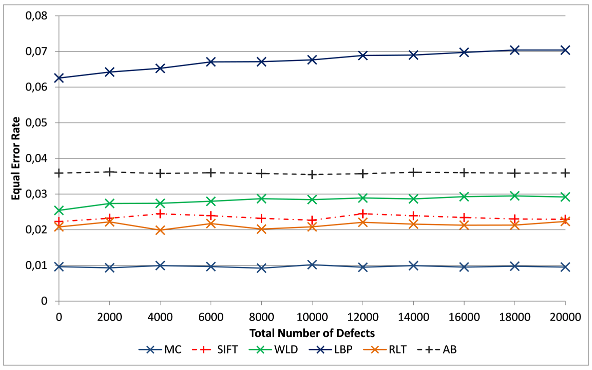

Fig. 43: EER hot and stuck pixels with denoising

9.2.3 Simulation Results with Denoising

While conducting the finger vein experiments all tests were run a second time. This time the denoising filter, which is used by LBP during preprocessing, was also applied for the other approaches. Table 6 shows that the baseline EER for MC, RLT and SIFT is slightly higher with denoising than without. It can be seen from figures 41, 42 and 43, that the impact of sensor ageing related pixel defects can be eliminated almost completely using denoising for MC, SIFT, RLT and AB, i.e. the EER remains almost constant even for a high number of pixel defects except some small statistical variations due to the random positions of the defective pixels. The EER of WLD and LBP rises only slightly. Although the baseline EER rises, the performance ranking for 0 defects is not changed. This ranking remains unchanged over the whole range from 0 to 10000 hot pixels, 0 to 10000 stuck pixels and 0 to 20000 combined hot and stuck pixels. Consequently, this indicates that denoising itself has an impact on the recognition performance but it is able to suppress the impact of sensor ageing related pixel defects almost completely.

Denoising can be made adaptive so that it is only used at a certain noise level and in addition only used for schemes where it is advantageous. Thus the baseline EER stays the same and if more pixel defects are present, the EER can be reduced.

9.2.4 Interpretation of the Results

MC MC tries to emphasise the centre lines of the veins, thus it generates binary images where the width of the vein lines should be 1 pixel. In addition MC uses a filtering operation to connect broken vein lines. Therefore it is quite insensitive to sensor ageing related pixel defects, which are point like, single pixel defects. On the one hand the filtering operation reconnects vein lines that might get broken due to defective pixels, on the other hand if a vein might get thicker due to stuck low pixels, this effect is mitigated by emphasising only the centre lines of the veins. If the number of defective pixels rises, vein lines might get broken because the filtering operation cannot connect them any more. This is why the performance of MC drops at some point. As it can be seen from the figures, MC is hardly influenced at all if only hot pixels are present. A hot pixel simply adds an offset to a pixel value, thus the pixel appears brighter. As the veins appear as dark lines in the image, bright pixels will only break the vein lines, which is corrected by the filtering operation of MC. Stuck pixels have a more severe influence on MC because they might be stuck at any pixel value between 0 and 255, i.e. also dark values. Thus, a stuck pixel can “extend” a vein line at least by one pixel and if there are several stuck pixels in the same image area, the filtering operation might connect the stuck pixel and “generate” a false vein line.

Using a denoising filter reduces the baseline performance of MC because it gets more difficult to accurately extract the centre lines of the veins as due to denoising the vein boundaries get less sharp than without it. But this simple denoising filter is able to mitigate the effect of hot and also stuck pixels on MC almost completely as the EER remains nearly constant not only up to 10000 hot or stuck pixels, respectively, but also up to 20000 pixel defects in total (10000 hot and 10000 stuck pixels).

SIFT SIFT tries to extract keypoints only out of “relevant” areas inside the image, i.e. it is also able to handle noise contained in the image. The keypoint extraction algorithm tries to extract the keypoints according to image information, like edges, high frequency content, etc. It is able to handle a certain amount of point like noise well. Like MC also SIFT is influenced more by stuck than by hot pixels. But the influence for a realistic number of pixel defects occurring in practical applications, which are much less than 1000 defects per sensor, is negligible. Even for 20000 defects in total its EER only rises from 0.02 to 0.032 which is a rise of 60%. Although SIFT is only the second best performing scheme, it is less influenced by pixel defects than the best performing approach, MC, which can be seen by comparing the percental rise of the EER. The influence of hot pixels might be lower than for stuck pixels because a hot pixel only adds an offset to a pixel value, i.e. the region containing the pixel stays smooth to a certain degree, whereas a stuck pixel has an arbitrary but fixed value, which might be totally different compared to its surrounding pixels.

Using denoising the baseline EER of SIFT gets higher but as with MC it remains almost constant up to 20000 pixel defects except some statistical variations. This means that also the SIFT based approach has no difficulties to deal with sensor ageing related pixel defects in practice.

RLT As it can be seen in figure 37, the vein lines do not get broken if Repeated Line Tracking is used as feature extraction method. But they appear brighter and have some additional noise on the vein boundaries. This is because of the way RLT works. It tries to extract the vein lines by repeatedly tracking dark lines inside the image. Dark means the pixel’s value is darker compared to its neighbours. If there are several hot pixels inside the image, a pixel which actually does not belong to a vein line might appear dark enough for RLT to count it as vein. RLT records the number of times each pixel is detected as belonging to a vein and applies a simple thresholding to these values afterwards to get the binary output image. This means that it is not independent of the vein width and especially if there are some hot pixels present along the boundaries of the veins, additional pixels, that are actually outside the vein but appear not as bright as the pixels further outside, may be detected as still belonging to the vein because the pixels further outside appear brighter due to the hot pixels. The EER of RLT rises from its baseline EER of 0.02 to 0.043 for 10000 hot pixels. Again the influence of stuck pixels is more severe as the EER rises to 0.051 for 10000 stuck pixels. The same as in the case of MC holds, i.e. a stuck low pixel is directly counted as vein line, thus especially stuck low pixels have a high influence on RLT. But also stuck high pixels lead to a drop in the recognition accuracy because they have the same effect as hot pixels and might lead to an increase in the width of the extracted vein lines with some additional noise on the vein boundaries.

The simple denoising filter is again able to mitigate the effects of sensor ageing almost completely. The EER of RLT remains quite constant even up to 10000 hot, 10000 stuck or 20000 combined hot and stuck pixel defects. Its baseline EER rises from 0.02 to 0.021 which is only a rise of 4%. Its EER for 10000 hot pixels is 0.0205, for 10000 stuck pixels it is 0.0211 and for 20000 pixel defects it is 0.0228 which results in a rise of 9% in the worst case. Thus also RLT can be made robust against sensor ageing related pixel defects using a simple denoising filter with only a small drop in the baseline recognition performance.

WLD WLD basically uses a simple local adaptive thresholding technique to extract the vein lines from the background. It compares the centre pixel to its neighbouring pixels inside a circular neighbourhood. If the centre pixel is influenced by a hot or a stuck pixel defect, this has a severe influence on the binary output image. As it can be seen in figure 37, the defective pixels are causing noise that appears as small circles in the binary output images. This is the main reason why WLD is heavily influenced by both, hot and stuck pixel defects. Again the influence of stuck pixels is higher than for hot pixels, that is mainly due to the rather low hot pixel offsets, which do not influence the adaptive thresholding so much. But it is the approach, which is influenced most by sensor ageing related pixel defects as its EER is more than 10 times higher than the baseline EER for 10000 hot and also for 10000 stuck pixel defects. This suggests that WLD is also sensitive to other types of image distortions, which may occur during practical use of finger vein recognition systems.

Again denoising is able to reduce the impact of the sensor ageing related pixel defects to a minimum. For WLD the baseline EER improves when using denoising as it is only 0.025 compared to 0.031 without denoising. The EER for 10000 hot pixels is 0.028, for 10000 stuck pixels it is 0.03 and for 20000 combined stuck and hot pixels it is 0.029, which is even lower than for 10000 stuck and only slightly higher than for 10000 hot pixel defects. In either case the rise of the EER is less than 17% (8.4% for hot, 16.3% for stuck and 14.8% for hot and stuck pixels combined).

LBP LBP compares the value of a centre pixel to its neighbouring pixels where the output value depends on the binary pattern of the neighbouring pixels inside a 3×3 neighbourhood. Here a multi-scale LBP operator was used, which repeats this comparison for several resized versions of the image trying to make the output more robust against noise and small perturbations. As stated above, LBP uses denoising during preprocessing by default. Thus the impact of hot and stuck pixels is relatively low compared to LBP without denoising which was also tested. This is quite obvious because already a single pixel, which changes its value inside the neighbourhood, changes the output value completely. Also the baseline EER of LBP would much higher without denoising and LBP was only evaluated with denoising. Using denoising the image is smoothed and the spiky noise is suppressed. Thus the impact of the pixel defects is manageable but LBP has the worst performance of all the tested approaches. Its baseline EER is 0.063, which is more than 6 times higher than the baseline EER of the best performing approach, MC. LBP is influenced slightly more by stuck pixels than by hot pixels only as its EER for 10000 stuck pixels rises to 0.069 compared to 0.067 for 10000 hot pixels only. Its EER for 20000 combined stuck and hot pixels rises to 0.07 which is a rise of 12% and thus the second highest drop in recognition performance after WLD (with denoising).

AB The adaptive binarisation or template matching approach is a generic adaptive thresholding technique, which is even simpler than WLD but it is much less influenced by sensor ageing than WLD. If the binary output images of AB are inspected, the finger boundaries can be clearly seen, whereas small veins do not appear in the output images. Single pixel defects mostly influence small structures inside the image like the small and thin veins but bigger structures like the finger boundaries are not influenced so much. Thus my assumption is that AB relies on the finger outline to a great extent and on thicker vein lines, which reduces its baseline EER compared to the other approaches, but makes it more robust against single pixel defects. For 10000 hot pixel defects and also for 10000 stuck pixel defects it performs better than RLT, LBP and WLD. For 20000 combined defects it performs nearly as good as MC. Its EER rises only slightly even without denoising. Its baseline EER is 0.036 and its EER for 10000 hot pixel defects is 0.043, for 10000 stuck pixel defects it is 0.051 and for 20000 combined defects it is 0.055. This means a rise of 17.5%, 40% and 50%, respectively. The fact that the rise in the EER is higher compared to hand veins shows that it does not solely rely on the finger outline as there is no finger or hand outline visible in the hand vein ROI pictures used during evaluation.

Using denoising the baseline EER of AB remains nearly the same (it actually slightly drops) and in addition it remains constant except some very small statistical variations for hot, stuck and also combined hot and stuck pixels. Thus in the case of AB denoising is able to mitigate the effect of sensor ageing on the recognition performance completely.

9.2.5 Finger Vein Conclusion

According to the estimation done with the empiric formula in section 8.4.1 a defect rate of 0.05 $defects/year$ would occur for the type of sensor used in the finger vein scanner. This defect rate is really low and would statistically lead to 1.5 defective pixels over a reasonable long sensor lifetime of 30 years. I showed that the impact of less than 1000 pixel defects per image is negligible for all of the tested approaches. Thus, the impact of sensor ageing and its related pixel defects for a reasonable, i.e. realistic number of defective pixels that will occur in practice, is negligible.

Even for the unrealistic number of 10000 (which is about 10000 times the number of defects the whole sensor would develop in 20 years) pixel defects, SIFT is hardly influenced at all because its key point extraction algorithm is robust against that type of noise. The influence on MC is quite low because it emphasises only the centre lines of the veins and uses a filtering operation, which reconnects broken vein lines and filters out some noise at the vein boundaries. RLT, LBP and the simple AB scheme are also relatively robust against the defects up to a rather high number of defective pixels. Only the performance of WLD drops significantly. Using a simple denoising filter the baseline performance slightly drops, i.e. the EER rises, but the influence of sensor ageing can be substantially reduced even for WLD. As mentioned above denoising can be made adaptive, so that it is only used if the performance drop is higher than it would be due to denoising.

In conclusion sensor ageing is not an issue for practical applications of finger vein recognition.

9.3 Hand Vein Results

This subsection presents the results regarding simulations and experiments done using the hand vein data. For evaluating the impact of sensor ageing on hand vein based biometric recognition systems, some of the feature extraction and matching schemes from finger vein recognition have been used. Specifically MC, SIFT, WLD and AB were used during the experiments. Again all simulations were run 5 times for hot pixels only, stuck pixels only and combined hot and stuck pixels using a defect rate of $4000\,defects/MP/year$ for 10 years with a time step of 1 year.

As the rotation correction done for finger vein images is based on the finger outline and cannot be done for the hand vein images (only the ROI images are used which do not contain the outline of

Method	Baseline EER	10000 Hot	EER (%)	10000 Stuck	EER (%)	10000 Hot+Stuck	EER (%)
MC	0.013	0.016	19.2%	0.017	25.8%	0.016	20.6%
MC *TA*	0.013	0.016	20.6%	0.022	61.4%	0.019	40%
SIFT	0.02	0.027	30.7%	0.03	47.6%	0.029	43.9%
SIFT *TA*	0.02	0.024	16.9%	0.023	14.1%	0.021	2.6%
WLD	0.073	0.112	52.9%	0.177	142.1%	0.161	120.2%
AB	0.157	0.163	3.9%	0.185	18%	0.174	11.2%

Tab. 7: Hand-vein EER baseline, 10000 hot, stuck and combined hot and stuck pixels

the hand or fingers), the rotation correction has to be done during the matching process. Therefore the Miura matcher was extended to repeat the matching process several times using rotated versions of the test image. Of course also the vertical and horizontal shift correction in the same way as with finger vein images is still performed (matching with slightly shifted versions of the image and using the highest matching score as result). I found out that the best compromise between runtime and matching performance is to use steps of 0.5° in a range of ±5°. A more sophisticated way of doing rotation and translation correction would be an alignment based on the vein crossings like it is done in [1]. On the one hand such an alignment is more complex and prone to errors because the vein patterns have to be extracted at first and then the vein "minutiae" have to be found and matched. On the other hand it will speed up the matching process and may improve the matching results. Note that this is not necessary for SIFT as it is rotation invariant. In addition, as the structures of hand veins are different compared, to finger veins the parameters of the preprocessing filters were adjusted. In particular the parameters of the Circular Gabor Filter and the High Frequency Emphasis Filter were slightly modified. Furthermore, the feature extraction parameters for WLD and AB were adjusted to account for the wider veins compared to finger veins.

9.3.1 Sample Aged Images

Figure 44 shows 2 hand vein images and the corresponding aged images containing 10000 defective pixels (5000 hot and 5000 stuck pixels).

9.3.2 Simulation Results

In table 7 the baseline EER for all the tested hand vein approaches and also the EER at 10000 hot pixels, at 10000 stuck pixels and at 10000 combined hot and stuck pixels is listed. Again also the percental rise of the EER compared to the corresponding baseline EER is given right next to each absolute EER value. By comparing the baseline EER values of the four tested approaches, it clearly shows that MC is the best approach again with an EER of 0.013. The performance of MC is lower for hand veins than for finger veins (baseline EER is more than doubled). The second best performing approach is SIFT, achieving an EER of 0.02 which is nearly the same as with finger veins and this time only 1.5 times higher than MC. On the third place is WLD with an EER of 0.073 which is much higher than for finger veins. The worst performing method is the simple thresholding scheme, AB with an EER of 0.157, which is also much higher than for finger veins. The other methods, i.e. RLT and LBP, were not tested because their performance was inferior compared to the the four tested schemes.

Figure 45 shows the performance in terms of the EER if only hot pixels are present in the probe images. First of all, it can be seen that the ranking of the 4 tested schemes does not change, even up to 10000 hot pixels. MC is the best, followed by SIFT, then WLD and the worst performing scheme is AB. It also shows that in the context of hand vein recognition MC is hardly influenced

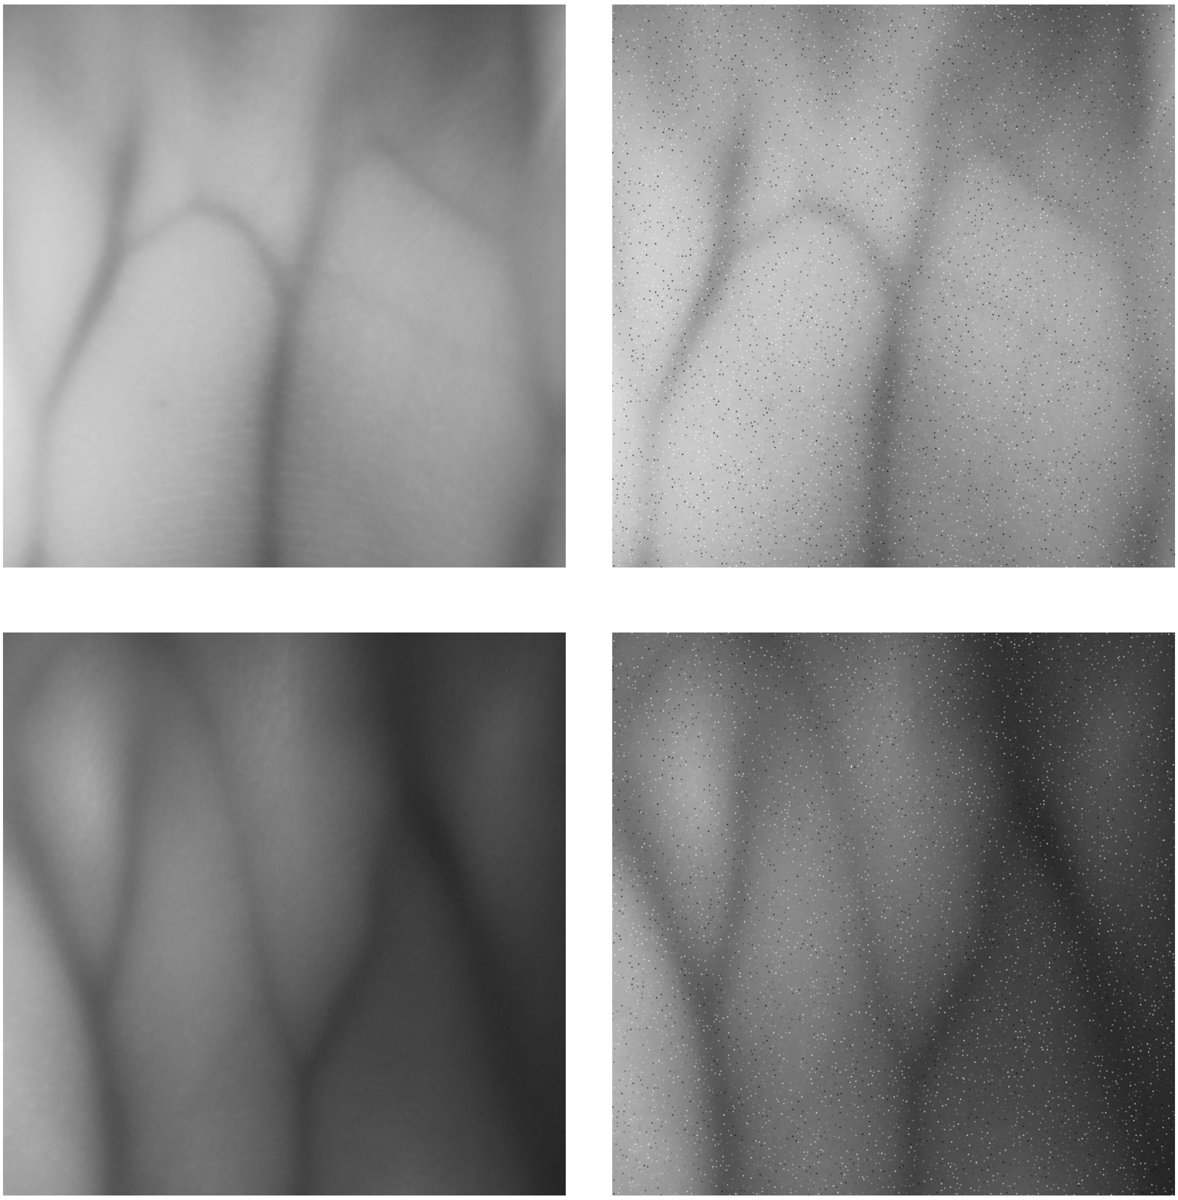

Fig. 44: Sample aged hand vein images, left: original image, right: aged image

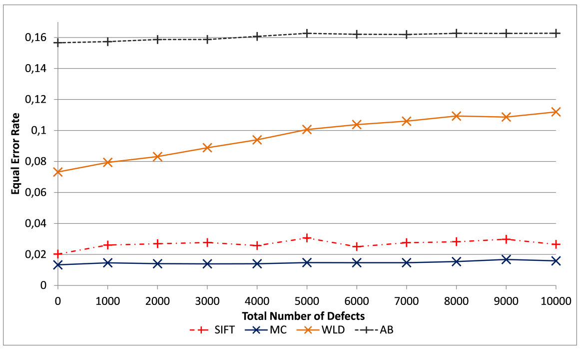

Fig. 45: EER hot pixels only

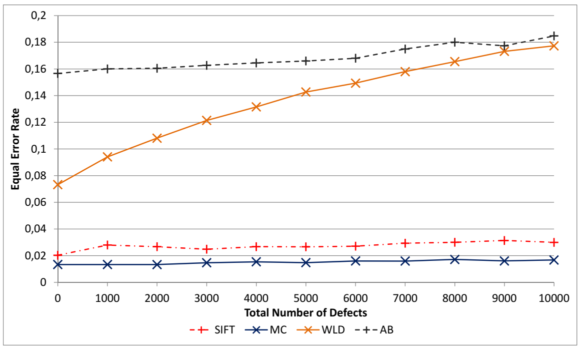

Fig. 46: EER stuck pixels only

at all by hot pixels, it EER at 10000 hot pixel defects is 0.016 which means a rise of 17.6%. Even SIFT is influenced more than MC for hand vein recognition as its EER at 10000 hot pixel defects is 0.027, which corresponds to a rise of 30.7% compared to its baseline EER. The simple AB scheme, although the worst performing approach, is influenced least. Its EER at 10000 hot pixels is 0.163, which corresponds to a rise of only 3.9%. Like for finger veins also for hand veins, WLD is influenced most by hot pixel defects and shows an EER of 0.11 or a rise in EER of 50.2%, respectively, at 10000 hot pixel defects. Due to the larger width of the hand veins compared to finger veins this rise is less severe than for finger veins.

The next figure, figure 46, shows the EER behaviour if the images are aged with stuck pixels only. Again the ranking of the 4 approaches does not change up to 10000 stuck pixels. The best approach is MC, followed by SIFT, then WLD and AB as last one. Similar to the situation with finger vein recognition also here the influence of stuck pixels is higher than of hot ones which is most visible for WLD and AB. The EER of WLD at 10000 stuck pixels has a value of 0.176 which is a rise of 140% compared to the baseline EER. This is nearly 3 times higher than for hot pixels. The EER of AB increases by 18% to 0.185 which is more than 4 times the increase for hot pixels. MC settles at an EER of 0.0168 for 10000 stuck pixels, which is equal to a rise of 26% and thus more than for hot pixels.

The following are the results for both, hot and stuck pixels combined, which can be seen in figure 47. Once again the ranking of all the tested schemes does not change during the whole range of 0 to 20000 defects in total, except that WLD gets worse than AB at 20000 defects. Here it is the same situation as with finger veins again. The impact of combined hot and stuck pixels is higher than for hot pixels only but it is lower than for stuck pixels only, which is another hint that stuck pixels have a higher influence in general. Please note that the x-coordinate ranges up to 20000 defects in total but to make a more meaningful comparison possible, I mainly state the values at 10000 defects in total. MC achieves an EER of 0.15 at 10000 defects in total, which is even better than for hot pixel defects only. SIFT settles at an EER of 0.029, which is higher than for hot but lower than for stuck pixels only. WLD has an EER of 0.162, which is a bit lower than for stuck pixels only and higher than for hot pixels only. AB is influenced less than for stuck pixels only and hits an EER of 0.174 for 10000 defects in total, again following the trend that the EER is higher than for hot pixels only and lower than for stuck pixels only.

9.3.3 Interpretation of the Results

MC First of all the performance of MC in the context of hand vein recognition is lower compared to finger vein recognition. This might be because there are less veins visible in the images compared

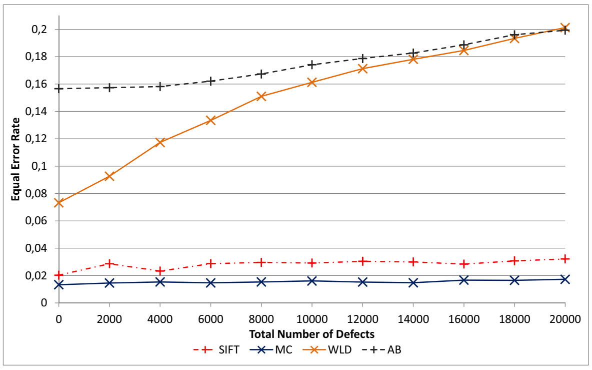

Fig. 47: EER hot and stuck pixels

to finger vein images, thus the extracted vein patterns exhibit fewer distinctions. As it can be seen in the figures 45, 46 and 47, the influence of the sensor ageing related pixel defects is lower compared to finger veins for all types of defects, i.e. hot pixels only, stuck pixels only and also hot and stuck pixels combined. Figure 44 shows some samples of aged hand vein images. It can be clearly seen that the vein structures are bigger compared to the relatively thin finger veins, i.e. hand veins are wider than finger veins. Thus the single pixel defects have less influence on the visible vein pattern inside the image, i.e. they cannot break the vein lines so easily. As mentioned before, MC only extracts the centre lines of the veins and uses a filtering operation to reconnect broken vein lines and to filter out noise at the vein boundaries. Due to that, the influence of single pixel defects on the feature extraction is rather low, even lower as with finger veins, as long as the number of defects stays within a reasonable range. Although the influence on MC for hand veins is lower than for finger veins, stuck pixels have a higher influence than hot and stuck pixels combined, which have a higher influence than hot pixels only like it is the case for finger veins.

SIFT The baseline performance of SIFT is nearly the same as with finger veins. Figures 45, 46 and 47 show that there are some variations in the EER values with an increasing number of sensor defects but there is still a clear trend visible. These variations are due to score calculation method used for SIFT, which is “count only” instead of “ratio score calculation” as it showed superior results. Count only means that the matching score is simply the total number of keypoints, which could have been matched between the two images. Ratio score calculation would divide this number by the total number of possible matches. Consequently, there are only discrete score values, thus the operating point for determining the EER cannot be set to such a fine degree (decimal values between two score values would not make sense). This also reduces the number of possible resulting EER values, which limits the accuracy of the resulting EER leading to the variations, that can be seen in the diagrams. In contrast to the situation for finger veins, when SIFT is used for hand vein recognition it is influenced more by the sensor ageing related pixel defects, i.e. the percental rise of the EER is higher than for MC. Hand veins exhibit bigger structures than finger veins, but the vein pattern has neither as many crossings nor intersections. In addition there is only a fewer number of veins visible in the image. This leads to a lower number of SIFT keypoints that can be extracted from the images. If there is the additional noise, induced by the pixel defects, present in the images, SIFT tries not to extract keypoints in noisy areas, which further reduces the number of keypoints that can be extracted. Consequently due to the count only score calculation, there are even fewer possible score values, i.e. less distinction between individual score values in general and between impostor and genuine scores in particular. The recognition performance drops because the score distributions of genuine and impostor scores start to overlap and thus the distinction between genuine and impostor scores gets more difficult (i.e. there are more false positives or more false negatives or both). Nevertheless, SIFT is still the second best performing approach.

WLD Like MC, the performance of WLD if used for hand veins is worse than for finger veins. Its baseline EER is more than twice as high compared to finger veins. This is mainly because the adaptive local thresholding is optimized for smaller structures, i.e. smaller vein width. Of course I tried to optimize the feature extraction parameters to get usable binary vein images. The most important parameters to tune are the radius of the local neighbourhood and the thresholds. If the radius is too small, the images get highly noisy as WLD also detects ridges and valleys on the skin surface as vein lines. Thus, the actual veins are covered by these smaller but more frequently appearing lines which are no veins at all. If the radius is too large, only the most prominent, i.e. widest and darkest, vein lines are detected and extracted. Consequently, the binary output images do only show a few thick vein lines and cannot exhibit many distinctions any more. Therefore a suitable compromise between the noise on the one hand and too few vein lines on the other hand has to be found. I experimented with the parameters for quite a while and a baseline EER of 0.073 was the best that could be achieved. Again like MC, the relative influence on the recognition performance is lower on hand veins than on finger veins using WLD, which is clearly indicated by the percental rise of the EER. The influence is less because the hot and stuck pixels due to sensor ageing do not complete break vein lines or generate artificial vein lines that might appear as vein lines in the binary output image. Due to the larger radius parameter used compared to finger veins, the defective pixels do not cause noise appearing as small circles in the feature extracted output images, which were the main reason for WLD's bad performance in finger vein recognition. Besides that, WLD is again influenced most by stuck pixels only, followed by hot and stuck pixels and least by hot pixels only.

AB Following the trend of MC and WLD also the performance of AB in the context of hand vein recognition is worse than for finger vein recognition. This is another hint that AB really relies on the finger outline, which is strikingly visible in the finger vein images, to a certain extend at least. The hand vein images do not exhibit the outline of the hand or the fingers as only a ROI, which is a region extracted from the centre of the hand surface, is used. Thus AB can only rely on the veins, that are not as clearly visible at the finger or hand boundaries and so the recognition performance drops. AB achieves a baseline performance, i.e. an EER of 0.157 which is more than 4 times higher than for finger veins. It is again the scheme, which is least influenced by hot and stuck pixels as it can be seen from the percental values in table 7. Although the basic principle of WLD and AB is quite similar, both are using an adaptive thresholding technique, AB uses a much larger window size compared to WLD. Thus a small number of defective pixels does not influence the thresholding result as much as it does for WLD. The same argument as with MC and WLD can also be used in the context of AB. As the hand veins are wider than finger veins, the single pixel defects influence the visible structures of the veins to a smaller extent. Thus also the drop in the recognition performance is less compared to finger veins.

9.3.4 Simulation Results with Templates Aged

For MC and SIFT the experiments were run a second time, again all simulations 5 times, with the mean of the EER as final result. This time not only the probe images but also the gallery images (or template images) were aged. A comparison of the matching performance for MC and SIFT with (denoted as *TA*) and without templates aged can be see in figures 48, 49 and 50. With templates aged, the same trend as without can be seen. Hot pixels alone have the least influence, followed by combined hot and stuck pixels and stuck pixels have the highest influence.

SIFT behaves as expected, i.e. if the templates are aged, the performance is increased. The diagrams also show that the EER of SIFT *TA* is lower than the one of SIFT across the whole range of defects. I say this is the expected behaviour. If the templates are also aged, the images get more similar again compared to the case where only the probe images are aged. If the images are more similar, the additional noise introduced by the pixel defects should have less influence on the matching results. I already mentioned that SIFT tries not to extract keypoints in noisy image areas. If only one of the images to compare is aged, there is a different number of keypoints, which can be extracted. The aged image will tend to have less keypoints than the unaged one. Such an unbalanced number of keypoints may cause unambiguous matches and makes matching keypoints more difficult. Thus the matching performance drops. If both images are aged, only the total number of potentially matchable keypoints decreases which leads to a lower number of possible score values as explained before. But it does not necessarily decrease the matching performance. Summing up SIFT performs better if not only the probe images but also the template images are aged.

Looking at MC the situation is different. First of all if only hot pixels are concerned, there is not much difference if the gallery images are aged or not. There are only some statistical variations, but the general trend remains the same and neither the EER of MC *TA* nor the one of MC is significantly better. For stuck pixels only the situation changes. With an increasing number of stuck pixel defects present, the gap between MC and MC *TA* gets bigger. This general trend is not surprising. The interesting thing is that MC outperforms MC *TA*. Although the images are more similar in the case of MC *TA*, the recognition performance drops as the EER rises. The explanation for this behaviour is not obvious at first sight. Having a closer look at the scores distribution for 10000 combined defects, shown in figure 51a, it can be seen that the impostor scores are shifted to the right towards the genuine scores. Also the genuine scores are slightly shifted to the right. This indicates that the images appear more similar in general (also images from different hands). Figure 51b shows the FAR and FRR for MC and MC *TA*, revealing that the FAR of MC *TA* is increased. Note that FAR is the same as FMR and FRR is the same as FNMR here, respectively. Thus the number of false positive matches increases for MC *TA* in comparison to MC. This is the main reason why MC *TA* performs worse than MC. Now the question is, why do the false positive matches increase? This can be explained having a closer look on the feature extraction results (binary vein images) and by having a look on how MC generates the binary images. If more and more stuck pixels are present, the simple filter operation used by MC is not able to reliably reconnect broken vein lines and filter out the outliers at the vein boundaries. In addition it tends to generate "artificial veins", i.e. vein lines visible in the binary output images where there are actually no veins in the image. This is mainly caused by the stuck low pixels starting from a certain number of stuck pixels. While the defective pixels are randomly distributed over the image, their position and parameters do not change for the same image sensor. One assumption for the simulations is that all images are taken with the same sensor, else evaluating sensor ageing effects would not make much sense either. Consequently, the pixel defects are the same, especially they are at the same spatial location for all images inside the database. If they lead to broken vein lines and artificial veins, these are located at exactly the same position in every feature space image. Matching is done using a simple correlation measure, which compares to feature space images and calculates the correlation based on their corresponding pixel values (0 or 1). Thus the images appear more similar, not only images of the same hand, but also images from different hands as they exhibit the same pattern of artificial and missing veins. If images from different hands appear more similar, this induces more false positive matches and thus the overall matching performance decreases. Finally looking at hot and stuck pixels combined the situation is similar to stuck pixels only. Starting at around 8000 pixel defects, the performance of MC *TA* gets significantly worse than MC. At 16000 defective pixels the performance starts to decrease rapidly (EER rises faster), at 18000 defective pixels it gets worse than SIFT *TA*.

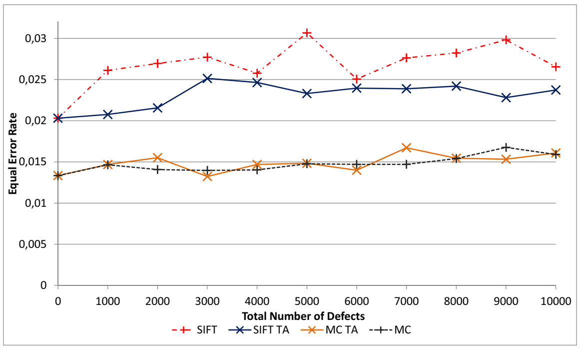

Fig. 48: EER MC + SIFT *TA* hot pixels only

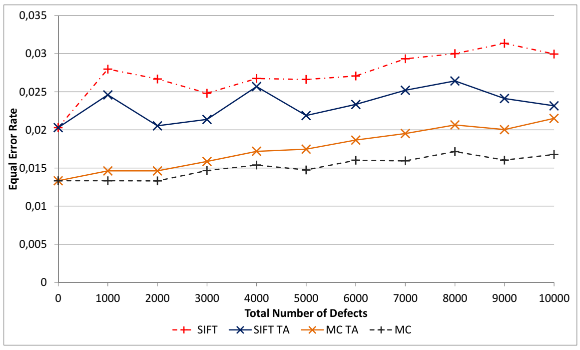

Fig. 49: EER MC + SIFT *TA* stuck pixels only

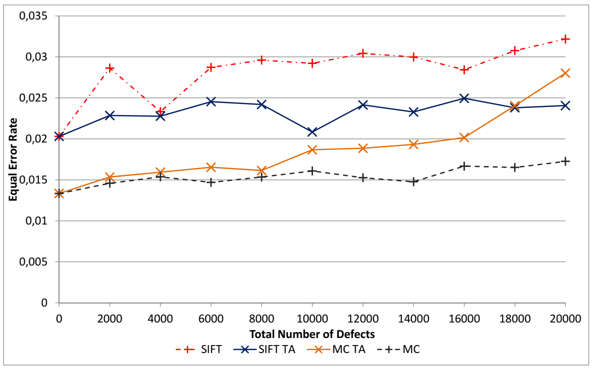

Fig. 50: EER MC + SIFT *TA* hot and stuck pixels

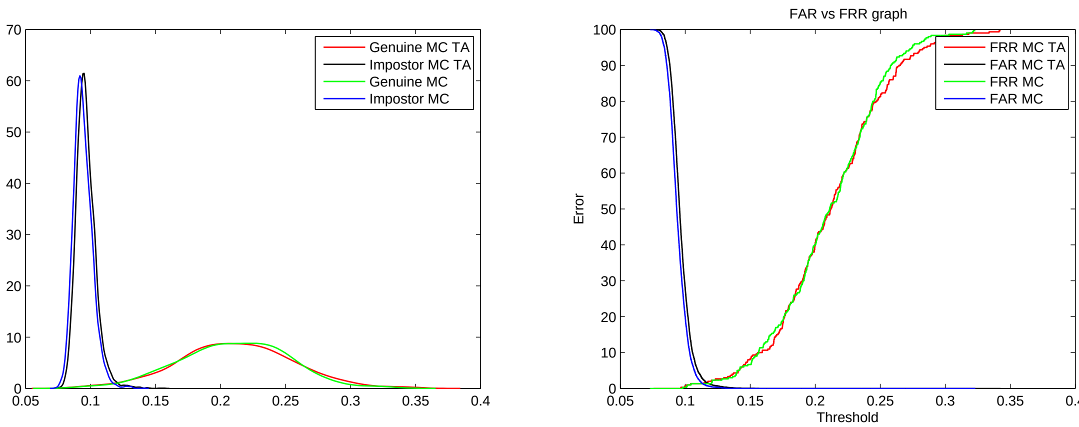

(a) Scores distribution

(b) FAR/FRR

Fig. 51: Scores distribution and FAR/FRR for MC and MC *TA* at 10000 defects

Note on the Fluctuations The attentive reader may have noted that although all simulations were run 5 times, there are still fluctuations in the EER values with an increasing number of defects, for some approaches more and less for others, but they are always present. Especially figures 48, 49 and 50 reveal this fluctuations because they only show a limited range on the y-axis, thus small variations become more apparent. First of all, I should explain why there are these fluctuations as they indicate that an increasing number of defects not always decreases the recognition performance. According to the pixel model, which was used as a basis for the simulations and thus for the experiments and evaluations, the defective pixels are randomly distributed over the sensor area. Consequently, they have random spatial positions in the images and also their parameters follow random distributions. However, their positions and parameters remain constant for all images taken with the same sensor at the same time, which was my assumption for all of the images in one database (either the ground truth or the aged ones). Now there are three things that can happen. First, a defective pixel (either a stuck or a hot one) may break a vein line at a specific location or may even extend a vein by at least one pixel if it is a stuck low pixel. With an increasing number of defective pixels these effect becomes more prominent, which leads to a general decrease of the recognition accuracy, because at some point the preprocessing and feature extraction methods are not able to cope with the increasing number of defects any more. This is of course also true for fingerprints, at which the ridges and valleys get interrupted by the defective pixels. This is what one would expected with an increasing number of defects. The second thing which may happen is exactly the opposite. A defective pixel may connect two parts of a vein line, where they appear broken due to bad image contrast, dust on the scanner or other distortions during the image capturing process. Actually it does not really need to reconnect the two parts, it may be sufficient if e.g. a stuck low pixel or a several stuck low pixels are located in the gap between the two vein parts. The preprocessing and feature extraction techniques are then able to reconnect the vein line, whereas without the additional stuck pixels they would not. Consequently. the matching performance increases. This is only one descriptive example. Of course this also applies for fingerprints and not only stuck pixels but also hot pixels can help to "improve" the vital parts of the image for feature extraction. Finally, a defective pixel may also have no impact at all, depending on its position. If it is a single stuck or hot pixel, located outside the veins or in the middle of a vein, the chance that it is filtered out during preprocessing and feature extraction is high and thus it has no impact on the matching result at all. Note that this effect is always per image only, i.e. a pixel defect which improves one image may corrupt another image. Hence an increase in the number of defective pixels does not necessarily lead to a decrease in matching performance. It may have no effect at all or even lead to an increase in matching performance. Returning to the initial question, this explains why there need not always be a linear increase in the EER with an increasing number of pixel defects considering a single simulation run. As a consequence, the simulations were run 5 times and using the mean of the EER as final result as this should smoothen the results. But statistically 5 runs are insufficient to cancel out all statistical variations due to the random positions and parameters of the defects. If the simulations are repeated 100 or 1000 times, the results will exhibit less variations. As a consequence of the high computational demand of most of the tested approaches and thus the high runtime of some of the tests, this was not feasible in the scope of this thesis.

9.3.5 Hand Vein Conclusion

As some of the approaches originally designed for finger veins were used during the hand vein evaluations, basically the same things that were concluded for finger veins also apply for hand veins. Although the baseline EER, except for SIFT, is higher due to the lower number of veins visible in the images, the influence of sensor ageing related pixel defects is less than for finger veins. Because hand veins are wider, single pixel defects do affect the visible vein patterns to a smaller extent and thus also the feature extraction is less influenced. Consequently, the drop in matching performance is relatively lower compared to finger veins. AB is influenced least, this time followed by MC and then SIFT. Again WLD is influenced most. According to section 8.4.1, a defect rate of $0.072\, defects/year$ would occur. Based on this defect rate, there would statistically occur 2.16 defects during a period of 30 years. It can be clearly seen that the influence of less than 1000 defective pixels is not worth mentioning. 2 defects would have practically no impact at all on the recognition performance. Although there is an impact for the high number of defects that were tested during the evaluations, is is still rather low for MC, SIFT and AB. Even though it was not tested, denoising could also be used in the context of hand vein recognition and is supposed to further reduce the impact of sensor ageing, most likely at the cost of a slightly increased baseline EER. The experiments with additional template ageing showed that SIFT behaves as expected as the results get better if also the templates are aged. MC behaves contrary, but only for stuck and combined hot and stuck pixels starting from a certain number of defects. A practical scenario where also the templates are affected by sensor ageing related pixel defects would be if the template database is created using an old scanner, which already shows ageing effects. In real life scenarios almost exclusively hot pixels occur, thus aged templates have no additional impact on hand or finger vein recognition systems.

In conclusion sensor ageing is not an issue when it comes to hand vein based biometric recognition systems.

9.4 Fingerprint Results

In the following, the results of the fingerprint experiments are given. As stated above, DB1 and DB2 from the FVC2004 data set were used during the experiments. FC, POC, NBIS and VF were evaluated. Due to the high runtime of FC and POC the simulation was only run once, for NBIS and VF it was run 5 times and the mean of the EER is used as final result. For NBIS and VF the experiments were repeated with aged gallery (template) images, denoted as *TA* in the diagrams and in the text.

MCC (Minutia Cylinder Code, see section 6.4.3 for details) was not evaluated. It is only a matching scheme. which relies on an external feature extractor. For this purpose the mindtct feature extractor from the NBIS package was used. The results of MCC were inferior to NBIS (EER more than twice as high as with NBIS). In addition also the runtime was several orders of magnitude higher compared to bozorth3 from NBIS. As mindtct is used for feature extraction and the sensor ageing related pixel defects mainly influence the feature extraction stage, evaluating MCC would mean to evaluate mindtct a second time with another matcher, which is the reason why I decided not to evaluate MCC.

9.4.1 Sample Aged Images and Minutiae Extraction

Figure 52 shows some sample images, the first and second row are from FVC2004 DB1, the third and last row are from FVC2004 DB2 (image 1_3 each time) and the corresponding minutia points extracted using VeriFinger. In the first and third row the unaged versions of the images are shown, second and last row show the aged versions with 10000 combined hot and stuck pixel defects present in the images.

As it can be seen for the DB1 image there are not many differences between the minutiae in the unaged and aged version. The aged image shows two additional minutia points in the top left and two in the middle left section (one above the seventh green line starting from the bottom and one left below the left triangle). All other minutiae in the centre part of the image do neither change nor disappear. Even the type and orientation of the extracted minutia points remain the same. This clearly indicates that minutiae based approaches are robust against sensor ageing related pixel defects as the extracted features are only faintly affected. If the feature extraction is robust against the defects, the matching process and thus the matching results will not change significantly.

The DB2 images show similar effects as the DB1 images. Most of the extracted minutiae points remain the same, also with the same type and orientation. Especially all minutiae located in the centre region of the image are not changed at all. Only some points on the right border of the image are missing in the aged version and in the top right area there are some additional points appearing. There is one interesting thing that can be seen on the left side of the image. The aged version shows two green lines, which indicate ridges, but there is not even a fingerprint visible at this position in the image. The same thing can be seen on the right edge of the image, where two additional vertical ridges appear (one on the top and one in the middle part). This indicates that presumably mainly the stuck pixel defects around the edges of the grey trapezoid in the image in combination with the preprocessing and filtering operations used by VeriFinger to enhance the image lead to false ridges and valleys. However there are no minutiae points extracted at these artificial ridge and valley lines. Thus it does not influence the feature extraction and the matching result. But this reveals that minutiae based approaches, at least VeriFinger, are not completely insensitive to sensor ageing related pixel defects and there might be a measurable influence for a higher number of defects.

9.4.2 Simulation Results DB1

The following table 8 shows the EER results for the FVC2004 DB1 achieved by FC, POC, NBIS and VF, which are the 4 matchers evaluated on fingerprints. First of all the baseline EER is given, followed by the EER values at a defect density of 40000 hot pixels per MP, 40000 stuck pixels per MP and 40000 combined hot and stuck pixels per MP, each one followed by the percental increase of the EER compared to the baseline one.

Note that for the fingerprint evaluations not the total number of defects present in the image is given. Instead the defect density per MP is shown. The images contained in DB1 and DB2 have different resolutions (DB1: 640×480, DB2: 326×364), thus giving a total number of defects per image would make comparing the results of DB1 against DB2 more difficult. Not the total number of defects per image is the crucial point, but rather the defect density, which corresponds to the influence on the recognition accuracy. The same total number of defects per image for DB1 and DB2 means, that the defect density is about 3 times higher for DB2 as the images have only $\frac{1}{3}$ the size of the DB1 images. Thus there are more defects per area and consequently the defective pixels are located closer to each other. The closer they are, the higher is the influence on the performance of the matchers as the defects then tend to break the ridge and valley lines or generate artificial ones. Consequently, the impact on the recognition performance if evaluating the DB2 images at the same number of defects per image is higher as with DB1. Using the defect density instead, a much more meaningful comparison between DB1 and DB2 results can be done.

Also note that the results for FC and POC exhibit a higher degree of fluctuations across the whole test range as the simulations were run only once due to the high computational demand and thus high runtimes of FC and POC. This is also the reason why the additional experiments with aged templates were only done with NBIS and VF and not with FC and POC.

As the baseline EER values show, the commercial matcher from Neurotechnology, VeriFinger, is by far the best performing matcher, achieving an EER of 0.025 for DB1. On the second place is FC with a baseline EER of 0.126, closely followed by NBIS, reaching an EER of 0.136 for DB1. The worst performing matcher is POC, which baseline EER is nearly 9 times higher then the EER of VF.

The impact of hot pixel defects is shown in figure 53. First of all, same as with hand vein recognition, the ranking of the tested fingerprint matchers does not change over the whole range from 0 defects up to a defect density of 40000 defects/MP. It can also be seen that none of the matchers is really affected, as the EER rises very slightly with an increasing defect rate. FC is influenced least, its EER actually slightly drops at 40000 defects/MP. FC does not exhibit many statistical variations across the whole range, i.e. the line is nearly a straight one, although the simulation and tests were

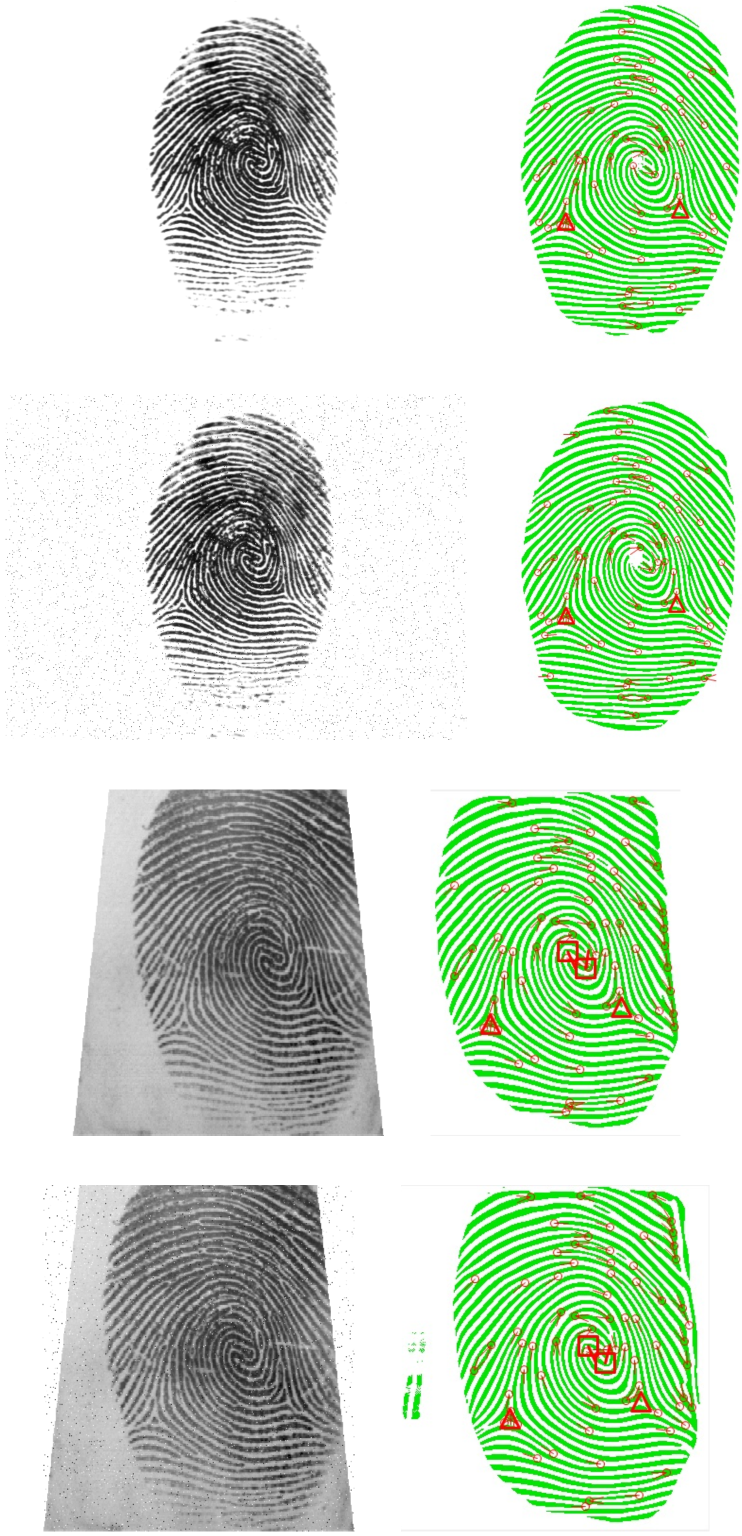

Fig. 52: Sample aged fingerprint images, left: original image, right: extracted minutia points

Method	Baseline EER	40000 Hot/MP	EER (%)	40000 Stuck/MP	EER (%)	40000 Hot+Stuck/MP	EER (%)
FC	0.126	0.125	-0.6%	0.13	2.8%	0.137	8.7%
POC	0.216	0.223	3.7%	0.204	-5.2%	0.199	-7.9%
NBIS	0.136	0.15	10.6%	0.156	14.9%	0.157	15.3%
NBIS *TA*	0.136	0.152	12.1%	0.153	12.9%	0.155	13.8%
VF	0.025	0.027	7.4%	0.027	10.6%	0.025	2.6%
VF *TA*	0.025	0.027	10.4%	0.029	16.5%	0.027	7.5%

Tab. 8: Fingerprint DB1 EER baseline and for hot, stuck and combined hot and stuck pixels

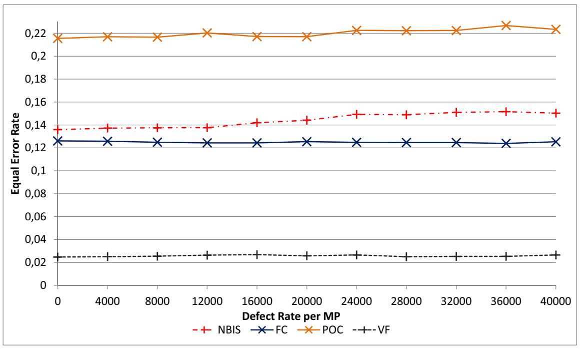

Fig. 53: Fingerprint DB1 EER hot pixels only

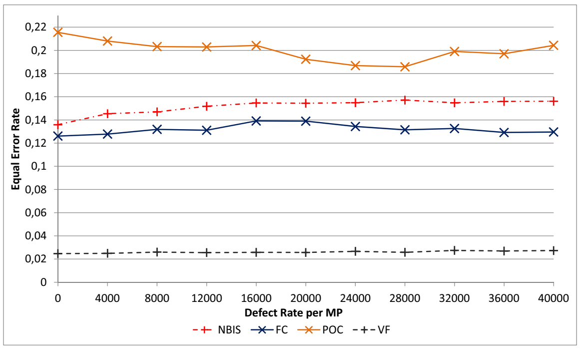

Fig. 54: Fingerprint DB1 EER stuck pixels only

only run once. This indicates that FC is robust against hot pixel defects to a great extent. The second robust matcher is POC, which EER rises by 3.7%, but the line shows some small variations. The two minutiae based matchers, NBIS and VF are influenced more than the ridge-feature based one and the correlation based one. NBIS is influenced most with an EER rise of 10.6%, Its EER remains quite stable up to 12000 defects/MP, then there is a rise in the EER and starting from 24000 defects/MP it remains almost stable again. VF is performing slightly better as its EER rises by 7.4% only and it rises almost monotonically. Nevertheless VF performs best of all tested matchers across the whole tested range in terms of EER and in addition it does not exhibit any major fluctuations regarding its matching performance.

Now the focus is on stuck pixels only. Figure 54 shows the impact of stuck pixels for DB1 on the 4 tested fingerprint matchers. Again the ranking of all 4 matchers does not change. VF is the best performing matcher, followed by FC, third is NBIS and the worst performing matcher is POC. NBIS shows a sharp rise in EER at the beginning, which continues up to 16000 defects/MP, but with flatter rise and then it remains almost constant at a value of about 0.155 until 40000 defects/MP. Looking at the percental rise values in table 8 reveals that NBIS is affected most by stuck pixel defects. Again VF is the best performing matcher, but it is affected second most if comparing the percental increases of its EER. The EER increases almost linearly with an increasing defect density. POC shows an interesting behaviour. At first its matching performance improves, i.e. its EER decreases up to 28000 defects/MP, where it starts to rise again. But its EER value at 40000 defects/MP is still below its baseline EER. This might indicate that POC performs better if stuck pixel defects are present, but it could also be an effect due to beneficial positions and values of the stuck pixels inside the image as explained in Paragraph 9.3.4. The EER of FC rises until 20000 defects/MP and then drops again. Nevertheless its value at 40000 defects/MP is above its baseline EER, i.e. an increase of 2.8%. I would

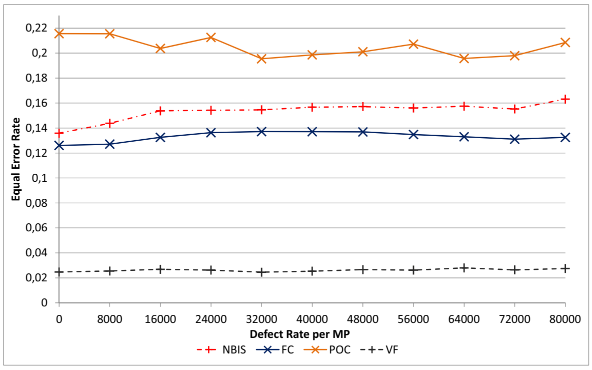

Fig. 55: Fingerprint DB1 EER hot and stuck pixels

argue that the rise at first followed by the drop is again due to statistical variations caused by the random positions of the stuck pixels and does not reflect a general trend. Once again, the minutiae based matchers are influenced to a greater extent than the non minutiae based ones.

The last figure 55 for DB1 shows the impact of combined hot and stuck pixels. Following the trend of hot and stuck pixels only, the ranking of the matchers does not change. VF again shows an almost linear increase in its EER, which is increased by 2.6% at 40000 defects/MP. NBIS shows a steep rise at the beginning up to 16000 defects/MP, then another slight rise up to 40000 defects/MP, and remains nearly stable up to 64000 defects/MP. Then there is a slight drop, followed by another steep rise starting from 72000 defects/MP. The EER of FC rises up to 32000 defects/MP, remains constant until 48000 defects/MP and then drops again. Its EER at 40000 defects/MP is 0.137 but at 80000 defects/MP its EER is only 0.13, which is slightly higher than the baseline EER. As the simulation for FC was only run once, this does not indicate a general trend. POC is again the worst performing matcher and shows the highest fluctuations. Regarding the percental increases of the EER values, this time VF is influenced least, which is quite interesting because it is influenced second most for hot only and also stuck only defects. NBIS is influenced most, which is in accordance with stuck and hot only defects.

Although there is of course an influence on the matching performance, this influence is not dramatically in any case. As the highest increase in EER is only about 15% even for an extremely high number of pixel defects, it is not worth mentioning in practice. In contrast to the situation when evaluation finger and hand vein recognition, here the influence of stuck pixels only is less then the influence of stuck and hot pixels combined. However hot pixels have the least influence again.

9.4.3 Simulation Results DB2

Table 9 summarizes the evaluation results for all tested fingerprint matchers on FVC2004 DB2. Comparing the baseline EER values reveals the performance ranking of the matchers. By far the best performing matcher is again VF, achieving an EER of 0.025 for DB2, same as with DB1. On the second place is NBIS with a baseline EER of 0.093. FC is able to achieve a baseline EER of 0.101, slightly better than POC with an EER of 0.104. Thus POC is the worst performing matcher, same as with DB1 and FC is on the third place for DB2, whereas it was the second best performing matcher on DB1. On DB2 both minutiae based matchers perform better than the correlation- and the ridge feature-based one. This indicates that not only the matching performance of each individual matcher depends on the data set, but some matchers perform relatively better compared to others, depending on the data set. The question why there are differences is not the topic of the present thesis, but an interesting question is if there are also differences in the sensitivity to sensor ageing related pixel

Method	Baseline EER	40000 Hot/MP	EER (%)	40000 Stuck/MP	EER (%)	40000 Hot+Stuck/MP	EER (%)
FC	0.101	0.103	2.5%	0.177	76.1%	0.155	54.1%
POC	0.104	0.104	0.5%	0.116	11.9%	0.111	7.6%
NBIS	0.093	0.099	6.5%	0.119	27.6%	0.109	17.3%
NBIS *TA*	0.093	0.109	16.8%	0.133	43.1%	0.119	27.4%
VF	0.025	0.025	-1.8%	0.029	16.5%	0.027	9.1%
VF *TA*	0.025	0.027	9.3%	0.03	20.4%	0.028	13.4%

Tab. 9: Fingerprint DB2 EER baseline and for hot, stuck and combined hot and stuck pixels

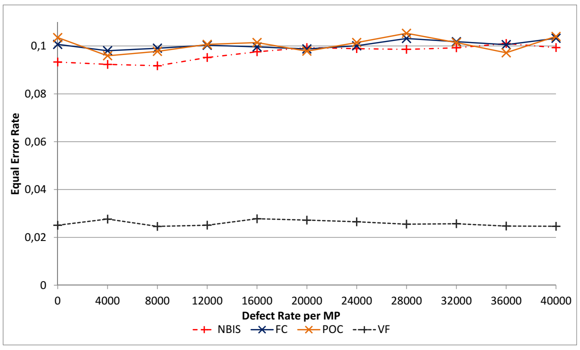

Fig. 56: Fingerprint DB2 EER hot pixels only

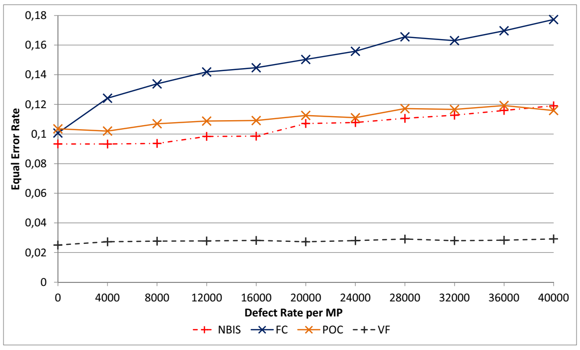

Fig. 57: Fingerprint DB2 EER stuck pixels only

defects between the two data sets. This question is going to be answered below.

The evaluations for hot pixel defects using DB2 show that first of all the recognition performances in terms of the EER of 3 of the tested matchers are close together. These are FC, POC and NBIS. Figure 56 shows that their performances stay close together if only hot pixels are present in the images. At first NBIS performs best of this triple up to a defect density of 20000 defects/MP. There all three perform almost equally with an EER of about 0.098. If the defect density is further increased, NBIS again performs best until 32000 defects/MP. Starting from about 34000 defects/MP, POC outperforms both, NBIS and FC. But at 40000 defects/MP, NBIS again performs best and FC and POC are having an almost equal EER of 0.104. As the fluctuations of the EER are rather large, no clear prediction which matcher is influenced most and which of the 3 is influenced least can be made. Also the percental rise values in table 9 can only provide some hints. Based on these, POC is influenced least, followed by FC and NBIS is influenced most. Up to now, only NBIS, FC and POC were compared against each other. Regarding the overall ranking, VF is influenced second most in terms of the percental EER change. But figure 56 clearly shows that the EER decreases starting at 16000 defects/MP up to 40000 defects/MP, where its EER is 1.8% lower compared to the baseline EER. Thus I would say VF is not influenced least but it is influenced least negatively by hot pixel defects. In addition, VF clearly shows a monotonically decreasing EER starting at 16000 defects/MP.

The situation for stuck pixels is different as it can be seen in figure 57. Here the ranking of the matchers does not change, except that POC outperforms NBIS at 40000 defects/MP and FC has a slightly better baseline performance than POC. Apart from that, the best performing matcher is by far VF, followed by NBIS. POC is able to achieve the third place and the worst performing matcher is FC. Regarding the influence of stuck pixels on the matching performance, this time a trend is more visible than for hot pixels. POC is influenced least as its EER increases just by 11.9% at 40000 defects/MP.

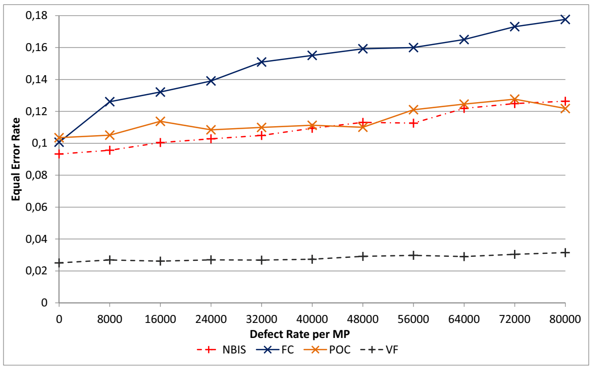

Fig. 58: Fingerprint DB2 EER hot and stuck pixels

Also VF is rather insensitive to stuck pixels, with an EER increase of 16.5%. Again VF shows a clear linear, monotonic increase in the EER towards higher defect densities. NBIS is influenced more than VF, but not as much as FC, which is influenced most. In contrast to the situation using DB1, here no clear trend towards minutiae based or non minutiae based matchers can be seen. It is also interesting that FC shows a steep increase starting from the beginning, which slows down at about 30000 defects/MP. On DB1 FC was quite stable and only a slight increase in the EER towards higher defect densities could be noted. This indicates that the impact of sensor ageing related pixel defects depends not only on the type of matcher, but also on the data set.

If both, hot and stuck pixels are present in the images, no clear decision if NBIS or POC performs better can be made. Nevertheless, figure 58 reveals that VF performs best across the whole range and FC performs worst. NBIS and POC are competing for the second place. This does not change the performance ranking up to 40000 defects/MP. From that point, POC performs slightly better than NBIS, but starting from about 50000 defects/MP the situation is changed once more until 76000 defects/MP, where POC outperforms NBIS again. Comparing the percental decrease in recognition performance or the increase in EER, respectively, it becomes clear that POC is again influenced least, followed by VF, then by NBIS and FC is influenced most. This is a completely different situation than it was with DB1. Once more, VF shows an almost linear increase in EER towards higher defect densities. The EER of FC increases sharply from 0 to 8000 defects/MP and continues to increase less steep afterwards.

Summing up the evaluation of DB2, POC is influenced least, by all of the tested defect types, i.e. hot pixels only, stuck pixels only and hot and stuck pixels combined (not taking the negative increase of VF for hot pixels into account). VF is influenced more than POC, but not as much as NBIS, thus it is on the second place considering the ranking. Consequently, NBIS is on the third place and FC on the last, i.e. it is influenced most. The evaluation of DB2 shows the same general trend that was found during the evaluation of finger- and hand vein recognition. Hot pixels cause the least decrease in recognition performance, followed by hot and stuck pixels combined, which cause a higher decrease and stuck pixels only cause the highest decrease in recognition performance. This is in contrast to DB1, where hot and stuck pixels combined lead to the most severe effects.

Comparison between DB1 and DB2 This paragraph outlines the differences between DB1 and DB2 images and tries to infer why the evaluation results of the fingerprint matchers are different. I will not go into detail why there are differences in the baseline performances, but try to explain the differences due to the impact of sensor ageing related pixel defects. First of all the images of DB1 and DB2 are different in image size. DB1 images have a size of 640×480 *pixels*, while DB2 images have a size of 328×364 *pixels*. Thus a defect density per MP was used instead of a total number of defects to compensate for the different image sizes. However, the visible fingerprint covers only a smaller percentage of the total image area in DB1 images compared to DB2 images. This can be seen in figure 32. The visible fingerprint covers about 70% of the image area in DB2 images and only about 30% in DB1 images. Another difference is that the background of the DB1 images is completely white (pixel value 255), whereas the background of the DB2 images varies from middle grey to light grey (pixel value 160 to 220) and is completely white only in the upper left and upper right outside image regions. Apart from that, other effects that may impact the recognition accuracy are present in both image databases, e.g. different and uneven pressure leading to different widths of the ridge and valley lines, interruptions across parts of the images (vertical and horizontal bright lines which break the ridge and valley lines), noise, etc.

The evaluations showed that in general the influence of pixel defects (hot and stuck) is higher for DB2 images than for DB1 images. The explanation for this general behaviour is quite simple. Although the visible fingerprint in the DB1 images covers less than half of the total image area compared to DB2 images, the absolute area covered by the visible fingerprint in pixels is still larger than for DB2 images. Consequently, the ridges and valleys are wider on average and thus the single pixel defects do have less effect on the ridge and valley pattern. This is the same situation as with hand vein images compared to finger vein ones. Now let us have a closer look on each defect type.

Hot pixels have a lower effect on the matching performance for DB1 than for DB2. This is also quite obvious as the background of the DB1 images is all white. A hot pixels adds an offset to a single pixel, but if the pixel is already white, it cannot get any brighter. Statistically the change for a hot pixel to be located in a valley is nearly the same for images of both databases. A hot pixel located in a valley, which has the colour of the background, will not affect DB1 images at all, but it might have a little effect on DB2 images as there the background gets brighter. If a hot pixel is located on a ridge, it affects images from both databases as ridges are dark and a hot pixel will cause a bright dot to appear inside a ridge.

Stuck pixels again have less influence on DB1 images than on DB2 ones. If it is a stuck high pixel, the same argument as with hot pixels holds. If it is a stuck low pixel, there is also a reason why the impact on DB1 is less than on DB2. The ridges appear darker in DB1 images than in DB2, where they are only dark grey, compared to nearly black for DB1 images. In combination with the pure white background this yields a higher contrast and a stuck low pixel located inside a valley can be more easily distinguished from ridge lines and thus filtered out during preprocessing or feature extraction. The same applies for stuck mid pixels.

Following the general trend, hot and stuck pixels combined have a higher influence on DB2 than on DB1 images, which is just a logical consequence of the situation for hot and stuck pixels only as it is explained above.

Another interesting observation is that on DB1 images, combined hot and stuck pixels have a higher influence than stuck pixels only, which is in contrast to the results for finger vein, hand vein and the DB2 evaluations. This observation cannot be explained using the higher contrast argument from above because then it would be exactly opposite. For VF the influence of hot and stuck pixels combined is even less then the influence of hot pixels alone. I cannot provide a conclusive explanation for this behaviour, but it might either be caused by statistical variations due to the small number of runs or due to the image enhancement techniques used by VF prior to the extraction of the minutiae points.

9.4.4 Interpretation of the Results

POC Summing up all results, POC was influenced least on DB2 and second least on DB1. As the matching performance on DB1 for stuck and hot+stuck pixels combined even increased, one could also say that POC was most insensitive on DB1, depending if sensitive means any change regardless in which direction or it only means negative changes. POC uses a fingerprint enhancement technique prior to feature extraction, which is designed to remove noise from the fingerprint images. This enhancement technique may also suppress most of the spiky shot noise caused by defective pixels. This is not surprising as both, POC and FC have been designed to be used especially for low-quality images and in the end, sensor ageing related pixel defects are just a kind of noise. POC correlates the images in the frequency domain which has some advantages like brightness invariance and immunity to noise. In addition, only a band limited version of the normalised cross spectrum is used, limited in a way that only those regions that are related to fingerprint data are contained in the spectrum. This additionally filters out interfering high frequency components like shot noise. Therefore POC is intrinsically robust against sensor ageing related pixel defects.

FC On DB1 FC was least influenced by sensor ageing related pixel defects compared to the other matchers, not taking into account that the performance of POC increased. On DB2 FC performed worst for stuck and combined hot and stuck pixels and is on the third place for hot pixels. FC uses a different feature extraction strategy than POC. Unlike POC, which uses the global ridge and valley structure of the entire fingerprint, FC uses this structure in a more local manner. Local ridge orientation and local ridge frequency are extracted and used during matching. Local structures are always more prone to small distortions like noise than global ones, thus FC is more sensitive to defective pixels. FC uses a normalisation step prior to feature extraction to adjust the images having a specific grey-level mean and variance. Especially stuck pixels have a high influence on the grey-level mean and variance in certain image regions, thus this normalisation step is of course affected. During feature extraction FC is based on a bank of Gabor filters, which are tuned to be responsive to the ridge and furrow structure of the fingerprint only. Therefore the feature extraction may again reduce the impact of the defective pixels. FC and POC both use the same additional fingerprint enhancement strategy prior to feature extraction. As FC is influenced more than POC, this enhancement strategy is able to reduce the impact of defective pixels only to a certain extent. Nevertheless, FC is in some cases more robust against sensor ageing effects than pure minutiae based matchers like NBIS, as it performed better on DB1 and on DB2 for hot pixels at least.

NBIS NBIS is a typical example of a minutiae based matcher. At first it generates a quality map to exclude image regions with low quality in the feature extraction step. Then a binarisation is applied to the images, mainly using morphological operations. This is followed by the minutiae extraction and an additional filtering of false minutiae. The remaining minutiae are then used during the matching process. Minutia points can also be regarded as local image features and thus they are affected more by pixel defects than global structures. This is confirmed by the evaluations as NBIS is the most affected matcher on DB1 and the second most affected one on DB2. NBIS does not use an additional image enhancement prior to binarisation. The defective pixels in combination with the morphological operations used during binarisation of the input fingerprint images may cause ridges to be interrupted and ridge ends to be connected to each other. As the sensor ageing related pixel defects only introduce shot noise, these image regions are not marked as bad quality in the first step and thus they are of course included in the feature extraction stage. The altered ridges and valleys cause some minutiae points to disappear, while others may be artificially generated. This can be seen on the DB2 image example in figure 52. These false minutia points may be filtered out, but some of them will remain after the filtering step, which influences the matching scores. There will be more minutia points that cannot be matched even for genuine matches, thus the genuine scores decrease, leading to more false negative matches.

VF VF is another minutiae based matcher. As it is a commercial one, there are no details about the principles behind the feature extraction and matching process available. The only thing I know is that it uses some advanced image enhancement techniques prior to feature extraction which is surely one of the reasons why it performed best during all the tests, on DB1 and also on DB2. On DB1 it is most robust against hot and stuck pixel defects combined and on DB2 its matching performance even increases if only hot pixels are present. This indicates that VF is really robust against hot pixel defects. According to Neurotechnology, VF is able to deal with sweat pores that are naturally present in fingerprint images exhibiting a certain resolution. Especially hot pixels located on a ridge are very similar to sweat pores and thus VF should be able to deal with hot pixels.

9.4.5 Simulation Results with Templates aged

The following figures 59, 60 and 61 and figures 62, 63 and 64 show the results with templates aged against only probe images aged for NBIS and VF on DB1 and DB2, respectively.

Let us first have a look on VF. Although the percental increases of VF's EER on DB1 suggest that VF *TA* is influenced more than VF, a closer look on the diagrams reveals that there is no clear trend visible. Sometimes VF outperforms VF *TA* and sometimes the situation changes, except for combined hot and stuck pixels starting from a defect density of 48000 defects/MP. There VF *TA* performs worse than VF, i.e. VF is influenced more by the defective pixels if also the template images are aged than if only the probe images are. But in general neither VF nor VF *TA* is influenced more compared to the other. Investigating the results on DB2 reveals that again VF *TA*'s percental rise of all the EER values is higher compared to VF, but again the diagrams show that there is no obvious trend visible. The two lines differ only slightly and again sometimes VF *TA* is influenced more than VF and vice versa, except for stuck pixels only where VF *TA* performs slightly worse than VF across the whole range. Thus I would say for VF there is no distinction if the templates are aged or not, because the matching performance does hardly change.

Now the focus is on NBIS. NBIS *TA* and NBIS perform equally for hot pixels on DB1. Although the percental rise of the EER at 40000 defects/MP may suggest that NBIS *TA* performs worse, the diagram clearly shows that there is absolutely no trend visible. However for stuck and combined hot+stuck pixels there is a trend visible. The diagrams and also the EER values show that NBIS *TA* performs better than NBIS, which is in a way the expected behaviour. If the templates are aged, the images will be more similar again and this should lead to improved matching results. Therefore NBIS *TA* behaves as expected on DB1. On DB2 the situation is completely different. For all three tested defect types, i.e. hot pixels only, stuck pixels only and combined hot+stuck pixels, the EER values in table 9 indicate that NBIS clearly outperforms NBIS *TA* with an increasing defect density. This is confirmed by the figures 62, 63 and 64. The question now is, why is NBIS *TA* influenced more than NBIS? Once again a look into the scores distribution can help to answer this question. The scores distribution for NBIS and NBIS *TA* at a defect rate of 24000 stuck pixel defects/MP is shown in figure 65a. It can be seen that again the impostor scores are shifted to the right towards the genuine ones. Figure 65b shows the FAR and FRR. There it can be seen that the FAR of NBIS *TA* is higher compared to NBIS. Both reveals an increased number of false positive matches for NBIS *TA* with an increasing defect density, which is the reason for the lower matching performance. Actually the same argument that was used to explain the expected behaviour can be used to explain this behaviour. On DB2 the lower contrast and darker background in combination with the stuck and hot pixels leads to more ridge and valley structures that are artificially generated and appear in the binarised images as it was explained in Paragraph 9.4.3. These structures appear in each image present in the database. If these lead to false minutiae during feature extraction, these minutia points will also be extracted from each of the images. Consequently, this leads to lower matching scores as these minutiae can always be matched. This lowers the genuine match scores but also the impostor scores and leads to

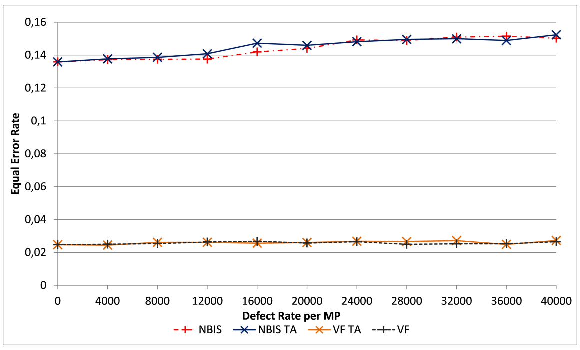

Fig. 59: Fingerprint DB1 EER NBIS + VF *TA* hot pixels only

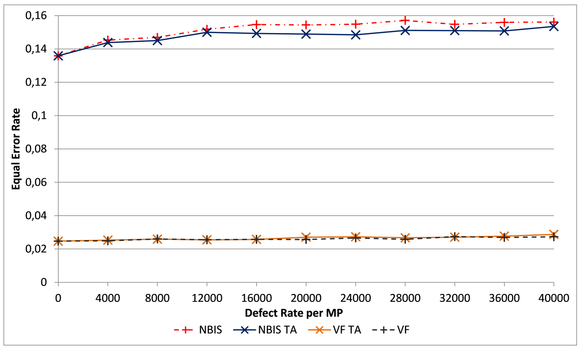

Fig. 60: Fingerprint DB1 EER NBIS + VF *TA* stuck pixels only

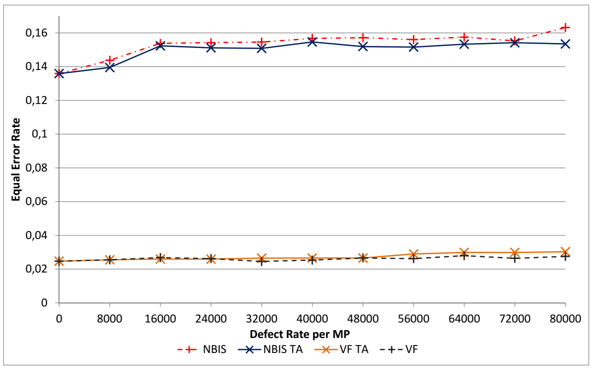

Fig. 61: Fingerprint DB1 EER NBIS + VF *TA* hot + stuck pixels

more overlapping of the genuine and impostor scores distributions. Thus the number of false positive matches increases and as a result the matching performance drops.

9.4.6 Fingerprint Conclusion

The evaluations on DB1 and DB2 showed that the impact of sensor ageing related pixel defects does not only depend on the type of pixel defects and on the particular matcher, but also on the images itself. The influence on DB1 images is in general less then on DB2 images. This is also true for the experiments with additional template ageing. On DB1 the behaviour is as expected, while on DB2 due to the higher influence of defective pixels in general it is not. On DB1 the minutiae based matchers are influenced more than the non minutiae based ones, which can be explained by the fact that minutia points are local image features that are more affected by single pixel defects. On DB1 FC is influenced least, followed by POC (if counting the positive influence just as influenced, else they would switch places), next is VF and NBIS is influenced most. The situation for DB2 is different. There POC is influenced least, followed by VF and NBIS, whereas FC is influenced most. Thus no clear trend, if minutiae based matchers are influenced more or not, can be seen. The only thing which can be seen is that VF performed best in terms of its absolute EER values for both databases. Although the influence on DB2 is higher, having a look at the two least influenced matchers, the rise in their EER still stays below 17% (for stuck pixels, it is even lower for hot ones). According to section 8.4.1 and depending on the fingerprint scanner used, a defect rate of 0.03 $defects/year$ (for the U.are.U 4000B

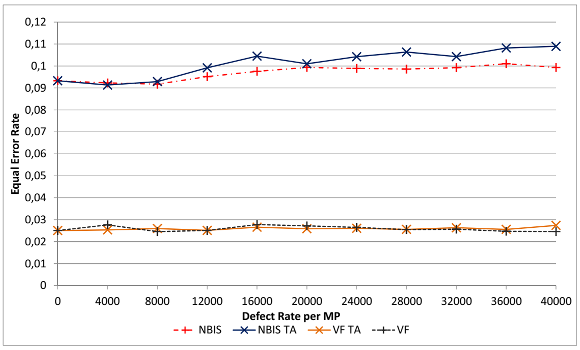

Fig. 62: Fingerprint DB2 EER NBIS + VF *TA* hot pixels only

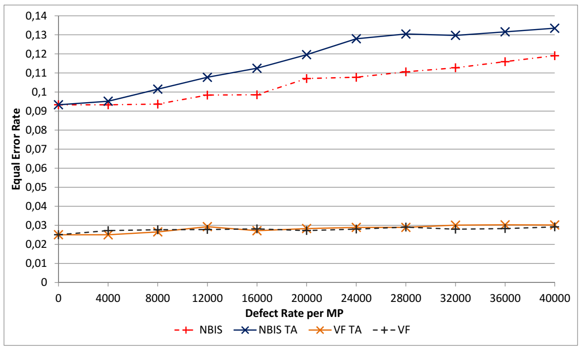

Fig. 63: Fingerprint DB2 EER NBIS + VF *TA* stuck pixels only

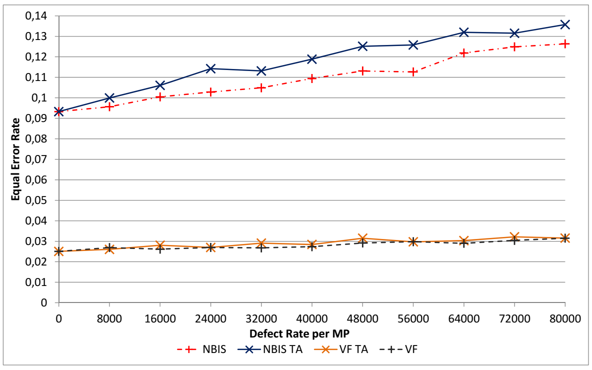

Fig. 64: Fingerprint DB2 EER NBIS + VF *TA* hot + stuck pixels

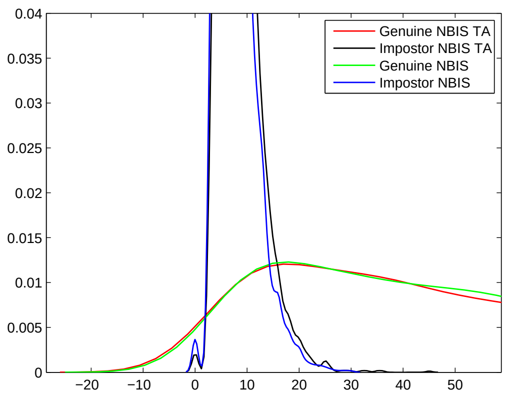

(a) Scores distribution

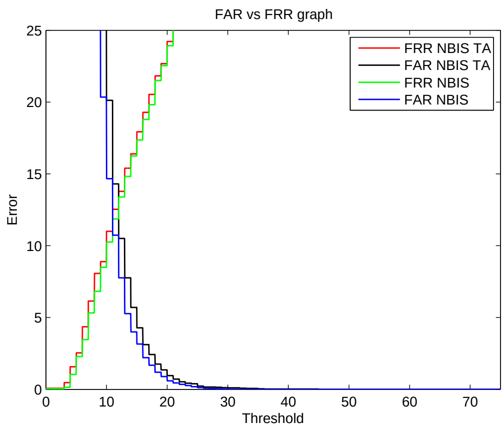

(b) FAR/FRR

Fig. 65: Scores distribution and FAR/FRR for DB2 NBIS and NBIS *TA* at 24000 stuck defects/MP

scanner) would occur. Statistically, this not even leads to a single defective pixel during a period of 30 years. Despite FC showing a steep rise in its EER starting from a low number of defects on DB2 for stuck and hot+stuck pixels combined, all other tested matchers only show a slight rise or even a decrease for a low number of defects. Stuck pixels are not reported to occur in real world applications. Furthermore, there would not be a single defective pixel and even if there are up to 100 pixel defects, the influence on the recognition accuracy is not worth mentioning. Following the trend of finger and hand vein recognition, sensor ageing is not an issue in fingerprint recognition systems.

10 Summary

The main research topic of the present thesis was:

"Evaluating the impact of sensor ageing on the recognition accuracy of finger image based biometric recognition systems".

According to my experimental results, there is no considerable influence on either of the tested schemes for a realistic number of defective pixels. In the following, I will give a short review of the present thesis at first and then I am going to conclude the evaluation results based on the different biometric traits, which are finger veins, hand veins and fingerprints to obtain this final statement. I will not go into details on each specific trait or any specific matcher once again as this was already done in the corresponding subsection of the results section.

This thesis started with a short introduction on biometric recognition systems, widely used for authentication purposes nowadays. Then the main impact of image sensor ageing, defective pixels, were described, followed by a quick overview over the existing literature regarding sensor ageing. Despite extensive literature research, by the time of writing this thesis, no literature could be found regarding the impact of sensor ageing on the recognition accuracy of biometric recognition systems. This lead to the main topic of the present thesis:

To be able to evaluate solely the impact of sensor ageing, the pixel defects had to be simulated, i.e. the images were artificially "aged". This procedure was selected in order to make sure that no other influences than sensor ageing related pixel defects are present in the images and thus affect the results. In the scope of the evaluations, fingerprint, finger vein and hand vein based feature extraction and matching schemes were evaluated: Two minutiae based fingerprint matcher, one ridge feature based matcher and one correlation based matcher for fingerprints, 5 binarisation type and one keypoint based approach for finger veins and 3 binarisation type and 1 keypoint based approach for hand veins. All evaluations were conducted using some well established data sets as ground truth and then ageing the images using the proposed ageing simulation algorithm. Usually only the probe images were aged, but in addition some experiments were conducted with aged template images one more time.

After the introductory section there was a short section about image sensors and their functional principles, to be able to explain the ageing related effects, which were described in the next section then. This section not only stated the different types of defects, it also mentioned some defect identification techniques and outlined the impact of different sensor parameters on the defect growth rate. Section 4 described several defect detection algorithms, which can be used to estimate the defect growth rate and the defect parameters from real world images. Section 5 explained the ageing simulation algorithm used to create the defect matrices and embed the pixel defects into the images.

Section 6 was about fingerprint recognition and also reviewed some general principles of biometric recognition systems. Besides fingerprint matchers also finger- and hand vein ones were evaluated, thus section 7 dealt with finger and hand vein recognition. Afterwards section 8 stated the experimental setup with the databases used for each biometric trait and the simulation settings which were estimated using an empiric formula. This was followed by a detailed analysis of the evaluation results, which can be found in section 9 and brings us finally back to this conclusion.

During the experiments, fingerprint, finger vein and hand vein matchers were evaluated and the results were analysed and interpreted. Each of the single results showed that for a realistic number of defective pixels, i.e. a realistic defect rate, there is no impact on the recognition accuracy. The estimated defect rates are between $0.03 - 0.07\, defects/year$ and would lead to $1 - 2$ defective pixels over a reasonable long sensor lifetime of 30 years. The evaluations showed that for such a low number of defective pixels none of the tested matchers was influenced.

There can be application scenarios where the defect rate is considerably higher. Pixel defects occur due to cosmic ray radiation, which is dependent on the altitude and and can lead to a pixel defect rate that is up to 300 times higher during transatlantic flights than on the ground. Such a scenario would be a biometric authentication system for air plane pilots. Then there could be up to 600 defects per image.

The evaluations showed that particular matchers are influenced by defective pixels due to sensor ageing, not only negatively but also positively. In general the influence is again negligible even up to 10000 defects per image for some of the tested schemes. This clearly shows that most feature extraction and matching schemes are quite robust against these single pixel defects, although they are not completely insensitive.

For finger veins a simple denoising filter was tested and the results showed that it is able to eliminate the impact of sensor ageing induced pixel defects almost completely at the cost of a slightly decreased baseline performance. If sensor ageing related pixel defects become a problem, such a simple denoising filter can be used to suppress the impact on the recognition performance. Although it was neither tested for hand vein recognition nor for fingerprints it is not far-fetched that denoising would further mitigate the influence of sensor ageing on these biometric recognition systems. If a simple denoising filter is really able to cancel out the influence of the defective pixels caused by sensor ageing on fingerprint and hand vein based recognition systems, will be part of my future work.

Referring back to the main question of this thesis I can now conclude that sensor ageing has no impact at all on real world applications of finger image based biometric recognition systems. Consequently, most state of the art fingerprint, finger- and hand vein recognition systems are already able to deal with the level of defective pixels caused by sensor ageing during every day operation, i.e. they are already robust against sensor ageing related pixel defects.

Yet another interesting research question would be if there is a noticeable impact due to sensor ageing related effects on non-optical biometric sensors, e.g. thermal fingerprint sensors. Therefore, at first the ageing processes on these sensors have to be studied. Afterwards a suitable defect model based on the ageing related effects has to be found, which can then be used in a simulation to evaluate the impact on the recognition accuracy like it was done for optical sensors during this work.

References

[1] Eui Chul Lee, Hyeon Chang Lee, and Kang Ryoung Park. Finger vein recognition using minutia-based alignment and local binary pattern-based feature extraction. *International Journal of Imaging Systems and Technology*, 19(3):179–186, 2009.

[2] Beining Huang, Yanggang Dai, Rongfeng Li, Darun Tang, and Wenxin Li. Finger-vein authentication based on wide line detector and pattern normalization. In *Pattern Recognition (ICPR), 2010 20th International Conference on*, pages 1269–1272. IEEE, 2010.

[3] Albert JP Theuwissen. Influence of terrestrial cosmic rays on the reliability of ccd image sensors-part 1: Experiments at room temperature. *Electron Devices, IEEE Transactions on*, 54(12):3260–3266, 2007.

[4] Albert JP Theuwissen. Influence of terrestrial cosmic rays on the reliability of ccd image sensors-part 2: Experiments at elevated temperature. *Electron Devices, IEEE Transactions on*, 55(9):2324–2328, 2008.

[5] Glenn H Chapman, Jenny Leung, Israel Koren, and Zahava Koren. Tradeoffs in imager design with respect to pixel defect rates. In *Defect and Fault Tolerance in VLSI Systems (DFT), 2010 IEEE 25th International Symposium on*, pages 231–239. IEEE, 2010.

[6] Glenn H Chapman, Jenny Leung, Ana Namburete, Israel Koren, and Zahava Koren. Predicting pixel defect rates based on image sensor parameters. In *Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), 2011 IEEE International Symposium on*, pages 408–416. IEEE, 2011.

[7] Glenn H Chapman, Israel Koren, and Zahava Koren. Do more camera pixels result in a better picture? In *On-Line Testing Symposium (IOLTS), 2012 IEEE 18th International*, pages 73–78. IEEE, 2012.

[8] Glenn H Chapman, Jenny Leung, Rohit Thomas, Ana Namburete, Zahava Koren, and Israel Koren. Projecting the rate of in-field pixel defects based on pixel size, sensor area, and iso. In *IS&T/SPIE Electronic Imaging*, pages 82980E–82980E. International Society for Optics and Photonics, 2012.

[9] Glenn H Chapman, Rohit Thomas, Zahava Koren, and Israel Koren. Empirical formula for rates of hot pixel defects based on pixel size, sensor area, and iso. In *IS&T/SPIE Electronic Imaging*, pages 86590C–86590C. International Society for Optics and Photonics, 2013.

[10] Jozsef Dudas, Cory Jung, Linda Wu, Glenn H Chapman, Israel Koren, and Zahava Koren. On-line mapping of in-field defects in image sensor arrays. In *Defect and Fault Tolerance in VLSI Systems, 2006. DFT'06. 21st IEEE International Symposium on*, pages 439–447. IEEE, 2006.

[11] Jozsef Dudas, Cory Jung, Glenn H Chapman, Zahava Koren, and Israel Koren. Robust detection of defects in imaging arrays. In *Electronic Imaging 2006*, pages 60590X–60590X. International Society for Optics and Photonics, 2006.

[12] Jozsef Dudas, Linda M Wu, Cory Jung, Glenn H Chapman, Zahava Koren, and Israel Koren. Identification of in-field defect development in digital image sensors. In *Electronic Imaging 2007*, pages 65020Y–65020Y. International Society for Optics and Photonics, 2007.

[13] Jenny Leung, Jozsef Dudas, Glenn H Chapman, Israel Koren, and Zahava Koren. Quantitative analysis of in-field defects in image sensor arrays. In *Defect and Fault-Tolerance in VLSI Systems, 2007. DFT'07. 22nd IEEE International Symposium on*, pages 526–534. IEEE, 2007.

[14] Jenny Leung, Glenn H Chapman, Israel Koren, and Zahava Koren. Automatic detection of in-field defect growth in image sensors. In *Defect and Fault Tolerance of VLSI Systems, 2008. DFTVS'08. IEEE International Symposium on*, pages 305–313. IEEE, 2008.

[15] Jenny Leung, Jozsef Dudas, Glenn H Chapman, Zahava Koren, and Israel Koren. Characterization of pixel defect development during digital imager lifetime. In *Electronic Imaging 2008*, pages 68160A–68160A. International Society for Optics and Photonics, 2008.

[16] Jenny Leung, Glenn H Chapman, Zahava Koren, and Israel Koren. Statistical identification and analysis of defect development in digital imagers. In *IS&T/SPIE Electronic Imaging*, pages 72500W–72500W. International Society for Optics and Photonics, 2009.

[17] Jenny Leung, Glenn H Chapman, Yong H Choi, Rohit Thomas, Zahava Koren, and Israel Koren. Analyzing the impact of iso on digital imager defects with an automated defect trace algorithm. In *Proc. SPIE*, volume 7536, page 75360F, 2010.

[18] Jessica Fridrich. Sensor defects in digital image forensic. In *Digital Image Forensics*, pages 179–218. Springer, 2013.

[19] Jenny Leung, Glenn H Chapman, Israel Koren, and Zahava Koren. Characterization of gain enhanced in-field defects in digital imagers. In *Defect and Fault Tolerance in VLSI Systems, 2009. DFT'09. 24th IEEE International Symposium on*, pages 155–163. IEEE, 2009.

[20] Jozsef Dudas. *Characterization and avoidance of in-field defects in solid-state image sensors.* PhD thesis, SIMON FRASER UNIVERSITY, 2008.

[21] Jenny Leung. *Measurement and analysis of defect development in digital imagers.* PhD thesis, Applied Science: School of Engineering Science, 2011.

[22] Thomas Bergmüller, Luca Debiasi, Andreas Uhl, and Zhenan Sun. Impact of sensor ageing on iris recognition. In *Proceedings of the IAPR/IEEE International Joint Conference on Biometrics (IJCB'14)*, 2014.

[23] Albert JP Theuwissen. Influence of terrestrial cosmic rays on the reliability of ccd image sensors. In *Electron Devices Meeting, 2005. IEDM Technical Digest. IEEE International*, pages 811–814. IEEE, 2005.

[24] Byoungjae Jin, Nohpill Park, Kayikkalthop M George, Minsu Choi, and Mark B Yeary. Modeling and analysis of soft-test/repair for ccd-based digital x-ray systems. *Instrumentation and Measurement, IEEE Transactions on*, 52(6):1713–1721, 2003.

[25] G.H. Chapman, I Koren, Z. Koren, J. Dudas, and C. Jung. On-line identification of faults in fault-tolerant imagers. In *Defect and Fault Tolerance in VLSI Systems, 2005. DFT 2005. 20th IEEE International Symposium on*, pages 149–157, Oct 2005.

[26] Davide Maltoni, Dario Maio, Anil K Jain, and Salil Prabhakar. *Handbook of fingerprint recognition.* springer, 2009.

[27] Pietro Coli. Vitality detection in personal authentication systems using fingerprints. 2008.

[28] Hiroshi NAKAJIMA, Koji KOBAYASHI, and Tatsuo HIGUCHI. A fingerprint matching algorithm using phase-only correlation. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 87(3):682–691, 2004.

[29] Lin Hong, Yifei Wan, and Anil Jain. Fingerprint image enhancement: algorithm and performance evaluation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 20(8):777–789, 1998.

[30] Anil K Jain, Salil Prabhakar, Lin Hong, and Sharath Pankanti. Filterbank-based fingerprint matching. *Image Processing, IEEE Transactions on*, 9(5):846–859, 2000.

[31] Raffaele Cappelli, Matteo Ferrara, and Davide Maltoni. Minutia cylinder-code: A new representation and matching technique for fingerprint recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(12):2128–2141, 2010.

[32] Barclays UK. Barclays newsroom, first in uk to lauch new biometric reader for customers. `http://www.newsroom.barclays.com/Press-releases/Barclays-first-in-UK-to-launch-new-Biometric-Reader-for-customers-bab.aspx`, June 2014. [Online; accessed 19-June-2014].

[33] B.T. Ton and R.N.J. Veldhuis. A high quality finger vascular pattern dataset collected using a custom designed capturing device. In *International Conference on Biometrics, ICB 2013*. IEEE, 2013.

[34] Jianjun Zhao, Hogliang Tian, Weixing Xu, and Xin Li. A new approach to hand vein image enhancement. In *Intelligent Computation Technology and Automation, 2009. ICICTA'09. Second International Conference on*, volume 1, pages 499–501. IEEE, 2009.

[35] Jing Zhang and Jinfeng Yang. Finger-vein image enhancement based on combination of gray-level grouping and circular gabor filter. In *Information Engineering and Computer Science, 2009. ICIECS 2009. International Conference on*, pages 1–4. IEEE, 2009.

[36] B. Ton . Vascular patern of the finger: biometric of the future? sensor design, data collection and performance verification, July 2012.

[37] Naoto Miura, Akio Nagasaka, and Takafumi Miyatake. Extraction of finger-vein patterns using maximum curvature points in image profiles. *IEICE transactions on information and systems*, 90(8):1185–1194, 2007.

[38] Naoto Miura, Akio Nagasaka, and Takafumi Miyatake. Feature extraction of finger-vein patterns based on repeated line tracking and its application to personal identification. *Machine Vision and Applications*, 15(4):194–203, 2004.

[39] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *Automatica*, 11(285-296):23–27, 1975.

[40] David 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.

[41] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *Computer Vision–ECCV 2006*, pages 404–417. Springer, 2006.

[42] Andreas Uhl and Y Holler. Iris-sensor authentication using camera prnu fingerprints. In *Biometrics (ICB), 2012 5th IAPR International Conference on*, pages 230–237. IEEE, 2012.

[43] Dario Maio, Davide Maltoni, Raffaele Cappelli, Jim L Wayman, and Anil K Jain. Fvc2004: Third fingerprint verification competition. In *Biometric Authentication*, pages 1–7. Springer, 2004.

[44] Jessica Fridrich. Digital image forensics. *Signal Processing Magazine, IEEE*, 26(2):26–37, 2009.

[45] Gayathri G Nampoothiri and Albert JP Theuwissen. Investigating the ageing effects on image sensors due to terrestrial cosmic radiation.

[46] Yap-Peng Tan and T. Acharya. A robust sequential approach for the detection of defective pixels in an image sensor. In *Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on*, volume 4, pages 2239–2242 vol.4, Mar 1999.

[47] Javier Alda et al. Bad pixel identification by means of principal components analysis. *Optical engineering*, 41(9):2152–2157, 2002.

[48] Christof Kauba, Jakob Reissig, and Andreas Uhl. Pre-processing cascades and fusion in finger vein recognition. In *Proceedings of the International Conference of the Biometrics Special Interest Group (BIOSIG'14)*, Darmstadt, Germany, September 2014.

[49] Karel Zuiderveld. Contrast limited adaptive histogram equalization. In *Graphics gems IV*, pages 474–485. Academic Press Professional, Inc., 1994.

[50] Olympus America Inc. Building a charge-coupled device. `http://olympusmicro.com/primer/java/digitalimaging/ccd/virtual2/index.html`, June 2014. [Online; accessed 19-June-2014].

[51] National Institute of Standards and Technology (NIST). Nist biometric image software (nbis). `http://fingerprint.nist.gov/NBIS/`, June 2014. [Online; accessed 19-June-2014].

[52] Neurotechnology. Verifinger 7.0 sdk. `http://www.neurotechnology.com/download.html#verifinger_sdk_trial`, June 2014. [Online; accessed 19-June-2014].

[53] Anouar Ben Khalifa, Sami Gazzah, and Najoua Essoukri BenAmara. Adaptive score normalization: A novel approach for multimodal biometric systems. *International Journal of Computer, Information Science and Engineering*, 7:18–26, 2013.

[54] EC Lee and KR Park. Restoration method of skin scattering blurred vein image for finger vein recognition. *Electronics letters*, 45(21):1074–1076, 2009.

[55] Jialiang Peng, Ning Wang, Ahmed A Abd El-Latif, Qiong Li, and Xiamu Niu. Finger-vein verification using gabor filter and sift feature matching. In *Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP), 2012 Eighth International Conference on*, pages 45–48. IEEE, 2012.