Continuous Biometric Authentication using Electrocardiographic (ECG) Data

Christian Pummer

MASTERARBEIT

eingereicht am
Fachhochschul-Masterstudiengang

Mobile Computing

in Hagenberg

im Dezember 2016

Declaration

I hereby declare and confirm that this thesis is entirely the result of my own original work. Where other sources of information have been used, they have been indicated as such and properly acknowledged. I further declare that this or similar work has not been submitted for credit elsewhere.

Hagenberg, December 5, 2016

Christian Pummer

Contents

Acknowledgments

The author wants to thank the department of Mobile Computing for enabling this thesis and the work it is based on. In particular, the author wants to thank Rainhard Dieter Findling for his guidance, mentorship and expertise throughout this thesis. Furthermore, the author wants to thank his parents, Herbert and Eveline Pummer for their steady support and efforts which only permitted this work.

This work has partially been carried out within the scope of *u'smile*, the Josef Ressel Center for User-Friendly Secure Mobile Environments, funded by the Christian Doppler Gesellschaft, A1 Telekom Austria AG, Drei-Banken-EDV GmbH, LG Nexera Business Solutions AG, NXP Semiconductors Austria GmbH, and Österreichische Staatsdruckerei GmbH.

Abstract

In the past decades, usage behavior and digital life style rapidly changed with emerging technologies such as smartphones, highspeed mobile telecommunication standards, social media, etc. While we are using countless digital services and devices on a regular basis, our main way of authentication remained unchanged: session-based authentication with tokens or secrets can be considered as de facto standard.

Continuous authentication might be a suitable concept to cope with those new conditions. While it would be impractical to continuously enter a password on a mobile phone, authentication just by touching the device seems tempting. Electrocardiographic (ECG) data can be continuously captured and verified. It is recorded by mere skin contact to ECG sensors.

In this thesis we design, build and evaluate a continuous ECG authentication system. Therefore, we record the FH Hagenberg Research ECG Database (FRED). We employ machine learning models for classification and finally evaluate system performance for identification and authentication use cases. Results indicate that continuous ECG authentication can achieve an equal error rate of about 7%. Unobtrusive data recording allows continuous ECG authentication to extend mobile device security, without necessarily reducing usability.

Chapter 1

Introduction

1.1 Motivation

"I think there is a world market for maybe five computers."

– allegedly Thomas J. Watson, IBM chairman, 1943

Obviously, this quote that is rumored to originate from Thomas J. Watson turned out to be underestimated. In the past decades, the number of computer systems grew rapidly and underwent a transformation from mainframe to truly personal computer. Smartphones, tablets, smartwatches, wearables and other personal devices contribute to the trend towards ubiquitous computing. Moreover, a new lifestyle emerged, where computer systems became more than tools dedicated for a certain purpose, but the real-world representation of the digital life. As devices like smartphones accompany us all day and serve us in every possible situation, they contain a lot of valuable information. From bank account, personal or business mail correspondence and health information, to private pictures, documents or social media accounts, our devices contain a lot of personal, sensitive and valuable information.

We try to protect our data from unauthorized access by employing authentication techniques. Traditional authentication schemes usually include secrets or tokens which are verified once per session, while security-wise it would be beneficial to perform authentication more frequently. This in turn demands for highly usable authentication technologies which require no or hardly any user interaction. While it is virtually impossible to continuously enter a 20-digit alphanumeric password, especially on mobile devices with limited user interface such as smartwatches, health trackers or wearables, many biometrics qualify for continuous use.

In search of new authentication technologies which qualify for continuous use while providing best possible security, the electrocardiogram (ECG) might be of interest. The ECG depicts the electrical potential of the heart over time and is commonly used for diagnostic investigation. It depends

on the individual physiology of the cardiovascular system, but cannot be derived by the phenotype of a person. The waveform of the signal has a characteristic shape, but varies between individuals. This individual variation in ECG signal can be exploited to distinguish people, to identify them by comparing their ECG to recordings within a database, or to test authenticity of a claimed identity. ECG is captured as easy as making skin contact by touching or approaching a sensor. In many cases, the interaction with a system already requires physical contact between user and system. Therefore, no additional interaction would be needed, to acquire a persons ECG. Furthermore, unobtrusive recording of ECG signals allows for the concept of continuous authentication, which is able to provide immediate intruder detection. Therefore, ECG is capable of adding security and usability to authentication systems.

1.2 Thesis Goal and Structure

In this thesis, we design, build and evaluate a system that continuously captures ECG data and provides authentication based on this data. In chapter 2, we study well known biometric authentication technologies, deal with their advantages and disadvantages and compare them with ECG authentication. Then, we review available technologies for capturing, processing and classifying ECG data. Chapter 3 covers basic concepts biometric authentication technologies commonly are built upon. The properties of ECG waves, recording and authentication technologies are explained in detail, as well as a comparative evaluation of existing ECG authentication approaches is presented in chapter 4. We proceed with designing and building our mobile, continuous ECG authentication system. We present our own approach in chapter 5. We record the FH Hagenberg Research ECG Database (FRED) using our hardware to evaluate and compare the classification performance of our system with state-of-the-art biometrics in chapter 6. Finally, we draw our conclusion in chapter 7.

Chapter 2

Biometric Authentication

Traditional authentication technologies provide access protection for a system in the beginning and for the duration of a session, i.e. the system is unlocked and accessible until it is locked by the user or automated mechanisms like timeouts or intruder detection. Authentication is provided based on something you know (knowledge-based) or something you have (token-based). These methods have several disadvantages. First of all, validity of the claimed identity is bound to a password or a key rather than to the user. If they are unavailable, e.g. a password cannot be remembered or a token was lost, authenticity of legitimate users cannot be verified. Furthermore, everyone who came into possession of those, e.g. by misappropriating or finding previously lost tokens or gaining access to passwords via phishing attacks, is able to authenticate on a system. To prevent unauthorized access to a system, e.g. by rainbow table or bruteforce attacks, long and complex passwords are required. In the case of password authentication, usability and security are contradictory requirements. The ever-growing number of devices and services, requiring long and complex passwords that shouldn't be reused across different systems makes the application of knowledge-based authentication increasingly inconvenient. For systems with limited user interface, where only a small (e.g. smartphone) or no keyboard at all (e.g. smartwatch, wearables, health trackers, etc.) is available, different authentication methods need to be employed.

2.1 Why Biometric Authentication

Biometrics are inherent behavioral or physiological properties of living beings. In other words, biometrics describe our appearance, *what* and *how* we are. Biometric traits include physiological properties like body height, face or fingerprint, as well as behavioral properties like voice, gait or signature. These properties are tightly bonded to *who* we are. The main advantages of biometric authentication, including usability, availability, security and

portability are directly inferred from this bond. Biometric authentication technologies are able to provide both security and usability while keeping the user interface minimal. They are usable, because other than passwords, biometrics have a small cognitive load. There is nothing to remember and therefore nothing to forget. For acquisition of many biometric traits, minimal user interaction is necessary, as stated in section 2.4. As authentication doesn't end in itself, but is a necessity to control access and use of resources, it shouldn't obstruct or hinder usage of those resources. Keeping the level of user interaction for authentication as low as possible, adds to usability and ultimately increases user acceptance. They are available, because most biometrics don't change significantly within weeks or months and they are on hand wherever we go. From some biometrics like voice, gait or ECG, even liveliness can be derived. They are secure, because many offer high classification rates, they are not prone to shoulder surfing and counterfeit attacks at least require a certain degree of equipment and preparation. And last but not least they are portable, because most biometrics can be measured with sensors that have small form factor and would easily fit into most mobile devices.

While the use of biometrics for authentication purposes offers many advantages, there are also security and privacy concerns arising. Biometrics contain personal and possibly sensitive information. System, service or authentication providers shouldn't gain access to biometrics, as the information contained in biometric signals could be abused or lost. Once the integrity of biometrics is compromised, they cannot easily be revoked or renewed. Therefore it is necessary to protect biometric patterns accordingly. One possibility is the application of biometric template protection schemes, that obfuscate original biometric information while maintaining authentication performance. Biometric template protection will be discussed in section 2.6.

2.2 Continuous Authentication

Regardless whether tokens, passwords or biometrics are used, authenticity is granted based on the assumption that the identity of the user remains the same for the lifetime of the session [24]. This assumption might have been justified in the past, but since computing became mobile and ubiquitous, it certainly no longer is.

For the means of security, it would be of benefit to repeat authentication as often as possible to ensure the authenticity of the user. Continuous authentication is the concept of performing user validation not only once in the beginning of a session, but repeatedly throughout the duration of user interaction. It is well known for using behavioral features, e.g. the interaction with the system like keystroke dynamics, mouse usage, application usage, etc. [17], but behavioral features might fail in providing sufficiently

high levels of security, causing biometric authentication to be limited to low security purposes and intruder detection. By combining continuous authentication with physiological biometrics, possibly higher levels of security can be achieved. However, while continuous authentication is beneficial to security, it shouldn't decrease usability and user acceptance at the risk of being deactivated.

2.3 User Friendly Authentication

Devices are routinely locked after periods of inactivity, which makes unlocking a quite frequent task in everyday life. Harbach et al. (2014) conducted an online survey and longitudinal field study regarding smartphone locking and unlocking behavior. The participants unlocked their phones on average 47.8 times per day, with a median of 42.1 and a standard deviation of 26.4 unlocks per day. Roughly a quarter of secured unlocks is perceived unnecessary by the users, with higher levels of dissatisfaction for private environments and secured unlocking mechanisms than for unsecured unlocks in public. As pointed out in [26], 57.1% of the participants do not use secured unlock mechanisms, with a majority stating "inconvenience" as the primary reason. Therefore it is inevitable for authentication technologies to adapt to the changed circumstances.

With traditional, high security authentication methods such as long and complex passwords, frequent authentication would increase inconvenience, lower usability and therefore decrease user acceptance at the risk of secured authentication being deactivated. System security can be improved, if users are continuously authenticated with a sufficiently strong authentication mechanism, while maintaining a high level of user acceptance and satisfaction. Luckily, the "what/who/how you are" type of keys, i.e. biometrics are permanently available and in some cases require no user interaction beyond regular use of the system.

The combination of continuous and biometric authentication can provide seamless use of a system, moving the task of user authentication away from the user, towards the system. Authenticity of users is checked continuously, providing immediate response to intruders. High levels of security can be achieved either by selecting highly secure biometrics, or by combining several highly usable biometrics into a multimodal framework.

2.4 Well Known Biometrics

There is a wide variety of different biometrics available for authentication today, as depicted in figure 2.1. This section provides an overview of several well known biometrics.

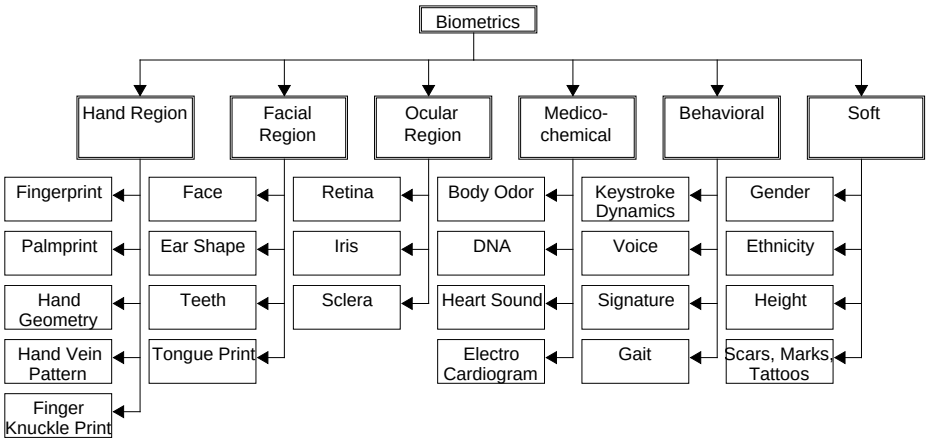

Figure 2.1: An overview of well known biometrics, structured by their body region or type [62].

2.4.1 Fingerprint

One of the most common biometrics used for authentication or identification are fingerprints. The discriminative nature of fingerprints has been known for more than a century. Usage is widespread and across several domains, from identification in forensic science, to authentication for access control on laptops. The ridges and valleys of the epidermis form a unique pattern of arches, loops and whorls. According to the study conducted by Unar, Seng, and Abbasi [62], most of the commercially available systems use features derived from characteristic points of the fingerprint, so called minutiae points. The traditional ink pad for fingerprint acquisition has been replaced by different sensors like optical, thermal, silicon or ultrasonic imaging sensors. Usage is widespread and generally accepted, e.g. in passports issued by the European Union since 2009 [84]. The sensors are small can be implemented easily into mobile devices. Availability of fingerprint biometrics can be considered as good, although it might not be applicable in some cases. While the amount of people suffering from adermatoglyphia – an extremely rare genetic disorder causing persons to have no fingerprints [69] – or without fingers are exceedingly small, regular conditions like wet, dry or unclean skin, scars, cuts, dead cells, skin or hand diseases, wrinkles, hard skin, etc. challenge fingerprint authentication. Data recording usually requires the finger to be placed on a sensor. Physical contact between skin and fingerprint sensor during acquisition might cause an oily or dirty sensor surface and distorted sensor readings.

Although fingerprints are easily acquired and counterfeited from surfaces that previously have been touched or high resolution cameras, fingerprints

are considered to offer a high level of security. They are used for identification and authentication on a personal, company and governmental level, since fingerprints are recorded and stored with the passport. Sensors are small and can be implemented in mobile devices. Current implementations are suitable for frequent authentication, but not for continuous, uninterrupted use. Systems are reported to achieve a false rejection rate (FRR) of 1-20% and a false acceptance rate (FAR) of 0.001-5% [10], depending on the desired purpose, data and system and evaluation setup.

2.4.2 Voice

Another prominent biometrics used for authentication is voice or language. It belongs to the group of behavioral biometrics, although the human vocal tract is a physiological feature. This is because not the vocal tract itself is compared for authentication, but the voice or language recorded over time. It is easily captured with microphones, prevalent in smartphones and other mobile devices. Usually, acoustic models are used to extract information and sound of speech samples. Features are derived either from the voiceprint of a subject, or from phonetic or phonological properties. While voiceprint, a plot of frequency over time with additional intensity information and phonetics, a branch of linguistics both are concerned with the sound of the human speech, phonology deals with the systematic organization of sounds in languages [71]. However, Bonastre et al. [8] mention that, although the name voiceprint suggests otherwise, it has nothing to do with a fingerprint. Other than fingerprints, human voice is subject to constant change. It changes with the time of the day, as well as with the time of the year or with age. It is also influenced by the speakers health or emotional state and can even be disguised on purpose. Although Unar, Seng, and Abbasi [62] attest voice biometrics an accuracy level of more than 90%, the performance of voice authentication certainly is a lot smaller compared to fingerprint or other biometrics. Voice biometrics are commonly used in multimodal authentication systems, where they can add additional features to the system and improve authentication performance. Very unobtrusive data recordings, inexpensive hardware and high user acceptance [16] make voice biometrics an excellent choice for extending existing frameworks.

2.4.3 Face

The face is the primary biometric characteristic used by humans to recognize each other [16] and one of the most versatile and powerful biometrics. In their survey, Abate et al. [1] state that face biometrics are the second most used biometrics after fingerprint and mention that, contrary to fingerprint, no consent is required for recognition tasks. The face can be recorded from a distance, e.g. by closed circuit television (CCTV), unnoticed and

without consent of the individual. Therefore, face recognition is not only used for authentication, but also for surveillance and tracking. This is also a possible threat for face authentication systems, as attackers could gain unauthorized access to pictures or 3D models of faces and conduct spoofing attacks. There is a wide variety of face recognition approaches, including 2D and 3D based systems, where data is acquired from still images or image sequences in visible light or infrared spectrum. Zhao et al. [66] divide face recognition systems into holistic approaches, and feature based approaches. While feature based approaches extract fiducial features like eyes, nose and mouth from the face, holistic approaches like eigenfaces [59, 61] extract features from the whole face, e.g. by performing transformations like principal component analysis (PCA) or discrete cosine transform (DCT). Most face recognition technologies use footage from regular cameras like smartphone or surveillance cameras. As hardware is already available in most off-the-shelf mobile phones, it is only a matter of software updates to globally distribute face authentication systems onto mobile devices. Over different approaches and datasets, an average accuracy level of 95% is reported in [62]. But authentication performance is highly dependent on the selected technology, use case, setup and data and conditions like illumination, facial expression or angle and the system's ability to cope with those changes. Usability of face authentication depends on the selected approach. For some approaches, no user interaction beyond regular use of a smartphone is required [1], others depend on moving a camera around the face to acquire a 180° pan shot [20] or challenge-response procedures such as pose or gaze estimation [2, 12, 22]. Face recognition can be considered as minimal invasive and socially accepted [16].

2.4.4 Eye

Together with retina and sclera vein patterns, iris patterns form the group of ocular biometrics. While they are located very close to each other, the recording procedure as well as their field of use are distinct from one another. As depicted in figure 2.2, sclera and iris lie on the outside of the eye and therefore can be recorded in the visible spectrum of light, e.g. with a smartphone camera. Recording is no more intrusive than taking a selfie. The retina lies on the inside of the eye, which makes recording more difficult. The eyeground needs to be illuminated with infrared light and the camera is located right before the eye. The close distance is determined by the diameter of the pupil and the area of the retina that should be recorded. Retina scanning is comparatively intrusive, needs special equipment and therefore usage is limited to military and similar purposes. The inconvenience and intrusiveness of data recording is disadvantage and strongest asset at the same time. It is very hard to gain unauthorized access to one's retina and nearly impossible to remain unnoticed in doing so. Images of iris or sclera

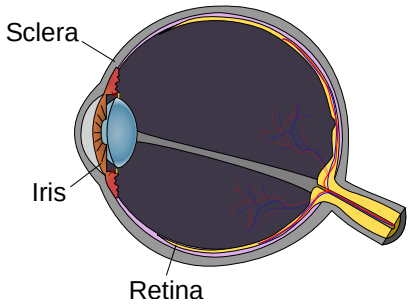

Figure 2.2: Illustration of a human eye, adapted from [80]. Sclera (the white part of they eye), retina (the light sensitive tissue in the rear of the eye) and iris (the colorful structure around the pupil) are commonly used for authentication.

can be obtained comparatively easy with high resolution cameras over a distance of several meters. Iris and sclera systems have to deal with different lighting conditions and reflections, but as long as the eye is clearly visible, even glasses or contact lenses do not negatively affect accuracy, according to Weaver [63]. Still, ocular biometrics are considered to be very secure, as they don't change considerably over time, offer a high level of uniqueness and accuracy levels of more than 99% [62].

2.4.5 Gait

Another biometric with advantageous properties is gait. It belongs to the group of behavioral biometrics and can be divided into approaches based on either acceleration or video recordings. Video based experimentation setups require cameras observing an area, where gait recognition should be performed. Lee and Grimson [41] present an approach using an orthogonal camera setup. As they rely on a static camera infrastructure, video based systems are not usable for mobile or continuous use and therefore mainly restricted to surveillance purposes. Accelerometer based systems are better suited for mobile and continuous use because they are more flexible and available, as most smartphones already have accelerometers included. Accelerometer based gait recognition is unobtrusive and socially accepted [16]. Muaaz and Mayrhofer [49] further divide accelerometer based systems into cycle based segmentation and fix length segmentation approaches. Cycle based approaches extract features from one or several cycles, while fix length approaches extract features from segments of fixed size. The extracted features are then subject to either template matching or stochastic/machine learning classification. Unar, Seng, and Abbasi [62] estimate the accuracy level of gait recognition to more than 90%, but mention that behavioral

attributes such as gait don't contain sufficient discriminatory information for reliable authentication. This is because they are affected by emotional state, health conditions, dietary habits and aging conditions. Additionally, according to Unar, Seng, and Abbasi [62] "identification based on behavioral biometrics is not beyond doubt since mimicking human behavior is easy for an experienced and skilled impostor". Nevertheless, behavioral biometrics can additionally add to security and usability of a system, e.g. as part of a multimodal biometric framework. Each additional biometric trait increases the key space and therefore increases the effort necessary for successful brute force attacks. Furthermore, the effort for spoofing attacks is increased, as additional biometric traits need to be mimicked. Security is enhanced, as additional biometrics add confidence in the result. User experience can be improved by lower response times due to short circuit evaluation, e.g. in identification use cases, behavioral biometrics can significantly reduce search space for potential matches.

2.4.6 Keystroke

A lot of research has been conducted in the field of keystroke dynamics. Similar to other behavioral biometrics, it offers high usability but low security. Keystroke dynamics recognition systems can be divided into systems that use static sentences and systems that perform authentication on dynamically changing text. Features derived from keystrokes include key hold and interval times, key press and release latencies [6] for static sentence approaches with physical keyboards. Systems that continuously capture user input from touch screens have even more features at their disposal, including the area occluded by the finger, pressure applied on the screen, gesture acceleration and velocity and many more [21]. For traditional, one time authentication for system access with a static sentence, as proposed by Balagani et al. [6], keystroke is considered not sufficiently secure [16, 45]. The gain in usability compared to password authentication is limited to not having to remember the password. In contrast, Bours [9] presented an approach that continuously captures text input by the user and unobtrusively performs user verification in the background. As soon as an intruder is detected by a change in keystroke dynamics, the system is locked and the attack is interrupted. The proposed system neither reduces usability, as keystrokes are captured continuously and no user interaction is required, nor security, as initial authentication is performed by another, more secure authentication technology. It can be considered as extension to existing authentication systems, unobtrusively improving security.

Bours [9] introduced another concept, called "concept of trust". Thus, negative authentication results don't instantly initiate the system to be locked. The system is only locked, if a trust level falls below a predefined threshold. Positive authentications add to the trust level, while negative au-

thentications decrease the level of trust. The threshold can be tuned to meet desired FAR and FRR. While keystroke dynamics relate more to desktop computers than to mobile devices, this principle is applicable to smartphones or other devices as well [4, 5, 23], although most likely different features need to be selected. For acquisition of keystroke dynamics and similar behavioral traits, no sensors are required, as the observed behavior is the interaction between user and system. Therefore, keystroke dynamics are highly unobtrusive user friendly and can be applied continuously.

2.4.7 Others

As shown in figure 2.1, there are many more biometric traits. In the following section, we mention some remarkable biometrics.

Signature

A fairly widespread, yet inconspicuous biometric is the handwritten signature. It has been used to sign documents, letters and contracts long before authenticity could be determined algorithmically. Baltzakis and Papamarkos [7] propose a system for signature recognition that is capable of correctly verifying 90% and identifying 80% of the signatures within a database of 115 individuals. The database contains a certain level of variation within each class, but no attempts were made to counterfeit signatures. Features include global statistics like height, width, number of edge points, cross points and closed loops, slant angle, as well as the grayscale values of the bitmap and statistical features regarding the occurrence of certain pixel sequences within a section of the signature image. While behavioral biometrics should not be trusted with positive authentication results, a negative authentication result could be reason for a more thorough examination of authenticity over different channels. The vulnerability to counterfeits might could be decreased using a real time requirement, but skilled impostors could still study and mimic the signature.

DNA

To the best of our knowledge, DNA is the most powerful biometric trait available. Except for identical twins, it is highly unlikely to find two individuals with the same DNA. Although people share about 99.9% of DNA, according to Korte et al. [37] the remaining 0.1% contain several megabytes of discriminatory information. Data is extracted from cells that contain DNA, such as skin, oral mucosa or the root of a hair. As it is impractical to examine the whole sequence of base pairs for recognition, most commonly *Short Tandem Repeats* (STR) are used. STRs are arrays of 5-50 *repeats* of 2-6 base pairs called *motif*. Those STRs are located on several 100,000 *loci*. About 20 loci are used in forensic science to identify individuals. In contrast to other

biometrics, familiar relationships can be derived from similarities in DNA. The system presented in [37] encodes the extracted features with an error correcting code and is claimed to reach an FAR as low as $m \cdot 2^{-73}$, where m is the number of enrolled individuals, and a FRR of 0.4%. As it is not possible to perform DNA recognition in real time, practical application is limited to forensics and criminalistic purposes.

Soft Biometrics

A very natural and intuitive way to distinguish people from one another is using soft biometrics. They include properties like gender, ethnicity, body height and distinguishing marks like scars, marks or tattoos, as shown in figure 2.1. Although a single property is unlikely to reliably distinguish individuals, a set of soft biometrics within a limited population has a certain discriminating power. Their significance and reliability are highly dependent on, selected features, feature probabilities and population size and therefore cannot be generalized. However, Reid et al. [54] mention that soft biometrics are obvious properties of individuals. They typically can be obtained from a distance and described by human understandable labels. Acquisition is nonintrusive and doesn't require consent. Therefore soft biometrics are commonly used for automated or manual surveillance and tracking purposes [54]. Furthermore, Unar, Seng, and Abbasi [62] state that soft biometrics can significantly improve performance of biometric authentication systems, especially in time constraint systems, by narrowing down the search space, e.g. if gender is taken into consideration in an identification application, about 50% of search space can be immediately removed. As soft biometrics are visual properties, they can be extracted from photographic or video footage.

ECG

Electrocardiogram is a well known technology for medical examination of the cardiovascular system and is explained in detail in chapter 4. It is the function of the electrical potential of the heart over time and is acquired as an electrical signal from the skin surface. The waveform depends on the individual physiology and has been shown to be a discriminative biometric property between people in many different studies [3, 18, 24, 29, 40, 44, 51, 56, 57]. Capturing ECG signals can be considered nonintrusive, as it can be integrated seamlessly in the normal use case of many systems, e.g. capturing the electrical signal while touching and holding a smartphone. Additionally, a continuous, uninterrupted ECG signal serves as liveliness and intruder detector. In section 4.4.4 some promising results documented in literature are presented.

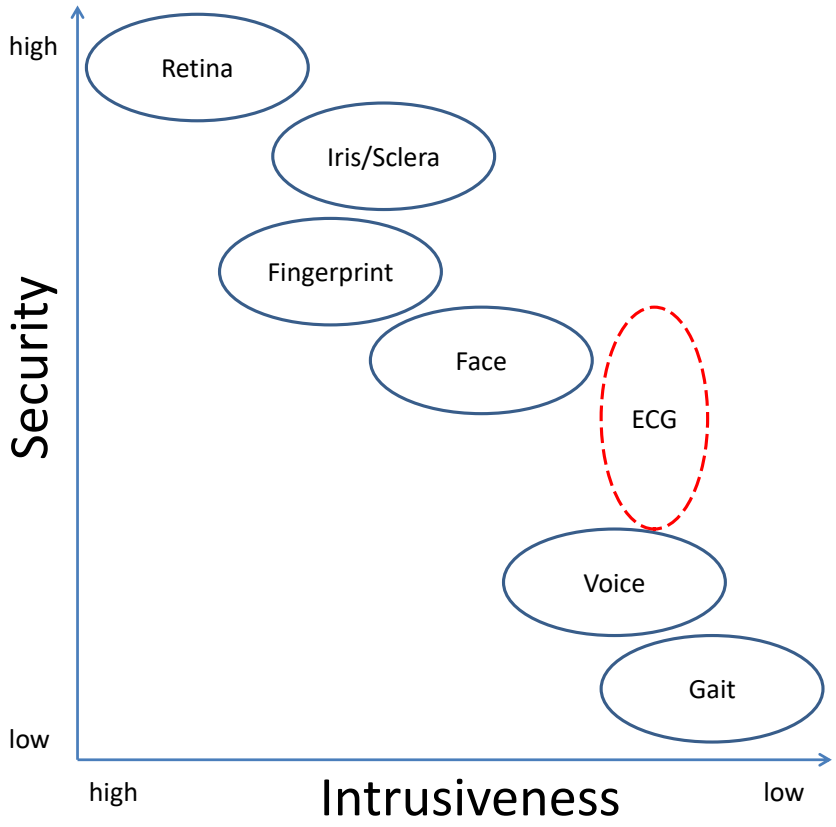

Figure 2.3: Illustration of estimated security and intrusiveness of common biometrics, based on [16, 52, 62].

2.5 Summary of Well Known Biometrics

The above mentioned authentication approaches use some of the most common behavioral and physiological biometrics. Each technology is a tradeoff between intrusiveness, fraud resistance and authentication performance. Figure 2.3 illustrates the estimated tradeoff and relations of intrusiveness and security of commonly known biometrics. The security estimation depicts the potential of biometrics, rather than the performance of specific approaches and is based on and combines the information presented in [16, 52, 62]. The level of intrusiveness of biometrics depends on the mode of interaction between user and system and can highly differ between different approaches within the same biometrics. We therefore ranked biometrics according to [45] and the best of our knowledge. Behavioral biometrics are thereby found to be the least intrusive, as typically very common behavioral traits are captured for classification. In the case of accelerometer based gait recognition, no explicit user interaction is needed at all, besides carrying a

device in the trouser pocket. Voice recognition requires rather intuitive, verbal interaction, but might not be applicable at all times or in all situations. Fingerprint recognition is judged to be about equally usable than face and iris/sclera recognition, but it might not be usable for continuous use. For many use cases of iris/sclera recognition, such as authentication on a mobile device, the intended use of the device already requires a mode of interaction between user and system, that allows for seamless integration of authentication processes. For example, we define the intended use case of smartphones to be held by users, looking at the screen from short distances of about 0.5 m at angles of approximately 90°. Within this mode of operation, the users iris/sclera could be continuously captured.

Physiological biometrics are by trend located on the more secure but less usable side if the chart. In general, security and usability of biometrics seem to be inversely correlated. When employing single biometrics for authentication, decisions have to be made between security and usability. Although ECG recognition is a relatively new idea, research conducted in the past decade [51] gives reason to believe that ECG recognition can provide usable yet secure biometric authentication. By increasing usability and security, ECG authentication can contribute to user acceptance of biometric authentication systems by eliminating "inconvenience" as the main reason for not using secured unlocking mechanisms [26]. However, most comparative studies and surveys do not include ECG authentication. Therefore, we used the results of the works presented in chapter 4 for our security estimate presented in figure 2.3.

2.6 Template Protection

Biometric authentication is able to positively influence digital security by making security usable, but simultaneously introduces some issues that need to be addressed. In order to evaluate authenticity of users, biometric templates need to be stored and compared to probes. Whenever data is stored, there is a chance unauthorized people gain access to this data. In the past, major companies have been subject to cyberattacks and hundreds of millions of datasets, in some cases containing personal information such as email, name or even passwords have been stolen [70]. It is strongly recommended not to reuse passwords on more than one system, program or website. When biometric authentication technologies are employed, we face several challenges. Breebaart et al. [11] state that unlike a PIN or password, our biometric characteristics are not renewable. If an attacker gains access to stored biometric features, spoofing attacks can be conducted.

Furthermore, as biometrics represent personal information, privacy concerns arise. One of the main benefits of biometric authentication - to rather link an account to a person than to a PIN - could turn out as a privacy

threat. Due to the distinctiveness of biometric features, attackers could link accounts across different, e.g. financial services and gain detailed insight into financial condition or investment plan of the customer. Besides, biometric data could contain sensitive, personal information like gender, age, ethnicity, etc., that could be used to acquire the identity of a person. Biometric data could also contain medical information and disclose medical conditions or risks for certain diseases, which could be used by health insurances to calculate insurance rate. Any leaked biometric data adds detailed, personal information to knowledge bases of countries, companies, intelligence services, criminals or other entities and can be used in any favorable or non-favorable way. Therefore special caution is advised when dealing with biometric information.

For the named security and privacy issues, it is advisable to prevent any conclusions from stored biometric data to identity or biometric characteristics of a person and to provide confidentiality, integrity and revoke- and renewability of biometric templates. In [11], the concept of Trusted Biometric Systems (TBS) is outlined and the need for an open standard for biometric template protection is emphasized and the meanwhile released standard for biometric information protection, ISO/IEC 24745:2011 [28] is referenced. When considering the use of biometric information within an application or system, it is advised to comply with this norm. Jain, Nandakumar, and Nagar [30] summarize and compare existing approaches for biometric template protection and conclude that there is no single best method for template protection, but that it depends on the use case and the selected biometric trait which approach serves best in securing biometric data.

While we won't go into further detail about biometric template protection, as it is outside the scope of this work, we acknowledge and emphasize the need for suitable template protection in every biometric authentication system.

Chapter 3

Building Blocks

Before we go into the details of ECG authentication in chapter 4, a short overview of tools frequently used for biometric recognition and authentication is provided in this chapter.

3.1 Distance Functions and Similarity Metrics

Distance functions are a tool commonly used for similarity metrics in template matching or to acquire difference vectors for classification. Some of the most common distance functions are presented below.

3.1.1 Euclidean Distance

Euclidean distance is one of the most common distance measures. In n-dimensional space, it is defined by the equation

$$d = \sqrt{\sum_{i=1}^{n}(q_i - p_i)^2}$$

as the metric distance between two vectors. Note that Euclidean distance is sensitive to scaling and shifting and emphasizes large values. For large vectors the distance is also large. Therefore normalization should be taken into consideration.

3.1.2 Manhattan Distance

Manhattan distance, also known as taxicab geometry is defined as the sum of absolute differences of their Cartesian coordinates. It is defined as

$$d = \sum_{i=1}^{n}|q_i - p_i|$$

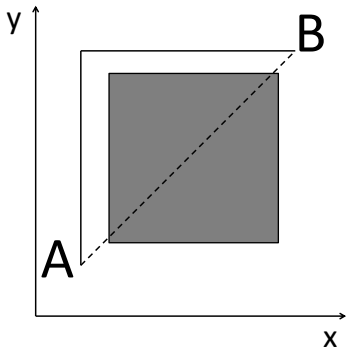

Figure 3.1: Minimal Manhattan distance between points A and B is depicted by the solid line. The dotted line shows minimal Euclidean distance. Minimal Manhattan distance is 2 units, while minimal Euclidean distance is $\sqrt{2}$ units.

in n-dimensional space. Figure 3.1 illustrates the minimal Manhattan distance between two points. It is similar to a taxicab driving in the streets of Manhattan that has to stick to the roads instead of crossing diagonally through the buildings.

3.1.3 Hamming Distance

Different to Manhattan and Euclidean distance, Hamming distance is not a geometric, but more a logical distance measure. It is defined as the number of positions, where the values of two vectors of equal length are different from each other. It is frequently used in coding theory for error detection and correction, but also as distance metric for classification, e.g. a sample x belongs to a class y, if less than z features of x differ from the class template of y.

3.1.4 Cosine Similarity

Cosine similarity depends rather on the orientation i.e. the angle between two vectors, than their magnitude and therefore is invariant to scaling. It is defined as

$$similarity_{ab} = cos(\theta_{ab}) = \frac{a \cdot b}{|a| \cdot |b|}$$

and evaluates to 1, if the vectors are parallel, 0 if they are orthogonal and -1 if they are diametrically opposed.

3.2 Classification Models

When class membership should be predicted based on previously made observations, classification models can be employed. After selection and extraction of problem dependent features, classification models are trained on existing data to learn relations between selected features and class membership to predict. For authentication, machine learning classifiers are employed to predict the probability of ECG probes to belong to the same user than ECG patterns stored in a database. Some of the most common classifiers used in ECG recognition are mentioned in this section.

3.2.1 k-Nearest Neighbors

The k-Nearest Neighbors (KNN) algorithm is used for classification as well as for regression. It is an instance-based, lazy learning algorithm, i.e. it doesn't make any generalization or abstraction, but it is constructed based on the samples of the training data, as explained in [39]. Classification of probes is based on class membership of samples - so called neighbors - which are located next to the probe in feature space. For regression, the predicted value for a sample is calculated by a summary statistic, such as mean or median of the k nearest neighbors. For the distance metric, most commonly Euclidean distance is used, but can be replaced by other distance metrics, such as Manhattan, Hamming, etc. As the output of KNN depends on a distance function, it is necessary to scale and center the features. Otherwise, features with large scales or off the center would contribute more to the result than features with small scales. If for some samples features are missing, the distance cannot be calculated. Incomplete features either can't be used or have to be interpolated.

When KNN is used for classification, the k closest samples in feature space are used to estimate the probability for class membership. If multiple classes have the same probability, the $k + 1^{th}$ neighbor is used or the class is selected randomly.

When tuning parameter k is selected very small, KNN can divide feature space into arbitrarily complex sections and is prone to overfitting. If k is selected too large, the boundaries may not represent the real complexity. As k is the only tuning parameter and k is a positive integer, a grid search for optimal tuning parameter is comparatively fast and easy and doesn't require expert knowledge on the classifier.

3.2.2 Nearest Center

Another common classifier used in literature is the nearest center or centroid classifier. Similarly to KNN, it is a distance based approach. The mean of all samples that belong to the same class is calculated for every feature. The

mean values of all features form a centroid. Each class is represented by only one centroid. A new sample is classified by assigning the class label of the centroid located next to it. Nearest center classifiers split the feature space linearly, which means they are only applicable, if the classes are linearly separable.

3.2.3 Linear Discriminant Analysis

LDA is a technique commonly used for dimensionality reduction and classification. According to Kuhn and Johnson [39], Fisher sought to find the combination of predictors that maximize the separation between classes in 1936. In 1939, Welsh was looking for a method to minimize the probability of misclassification. Eventually, both ended up in the same conclusion. LDA finds a linear hyperplane, that optimally separates two classes from each other. Additionally, it can provide valuable information of the relative importance of predictors, when looking at the linear discriminant function coefficients.

While LDA is capable of finding an optimal solution, it's performance strongly depends on two preconditions. The number of samples has to be larger than the number of predictors. In [39], the ratio is specified by at least 5 to 10. It should be noticed, that resampling techniques like cross validation reduce the sample size by a certain factor. For 10-fold cross validation, the number of samples should therefore exceed the number of predictors by at least a factor of 50 to 100. Besides, the predictors should be uncorrelated. For best performance, it is a common practice to use PCA to receive uncorrelated predictors before applying LDA. It is recommended to center and scale the data. Once these requirements are met, LDA is an easy to use linear classifier.

3.2.4 Neural Networks

Neural networks (NN) are powerful, nonlinear (for NNs consisting of more than one neuron) algorithms for both regression and classification. Their design is inspired by the very neurons in the human brain. While the human brain consists of about 10^{12} neurons, the simplest artificial neural network is called "perceptron" and consists of only one neuron. Each artificial neuron consists of components to receive information from other neurons, process this information and transmit it to other neurons. Although the theory of NN dates back to 1943 [47], they only became relevant with the rise in computational power during the past decade.

NNs belong to the class of supervised learning algorithms. When a perceptron is trained for regression, an input vector is passed to the neuron. The information is processed by multiplying each element of the input vector with a corresponding weight from a weight vector. The weights are initialized randomly. The weighted input signals are then summed up. The resulting

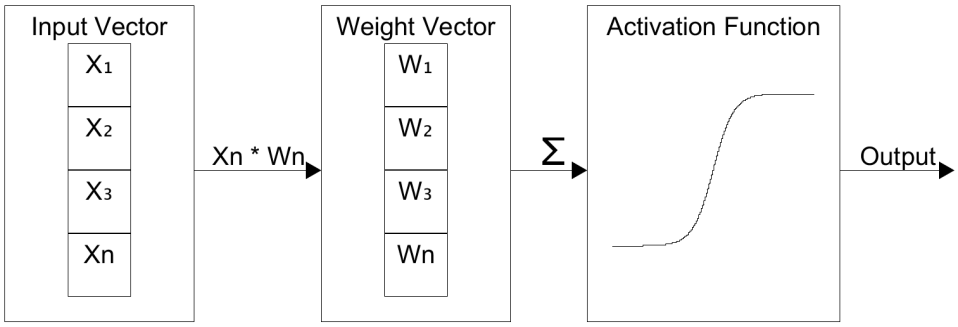

Figure 3.2: The main components of an artificial neuron.

value is then passed to an activation function, commonly a sigmoid or rectified linear unit function. During the learning process, the predicted number is compared to the real value. Depending on the error, the weights are tuned until predicted and reference value match. In more complex NN topologies consisting of more than one neuron, typically arranged in layers, the error is propagated to all neurons and the weights are tuned accordingly.

If NNs are used for classification, the topology diverges slightly. The hidden layers still calculate continuous values from inputs and weights, but after the last layer, an additional output function (e.g. sigmoidal) is used. For every class, there is one output neuron that outputs a value between 0 and 1, which must not be mistaken as class probability, as the values do not add to 1.

NNs are rather complex algorithms with several tuning parameters. Their complexity and power rises with the number of layers in a network and neurons per layer. Large networks are able to perfectly adjust and give optimal results for training data, but therefore are prone to overfitting. Tuning parameters include network topology, numbers of layers, neurons per layer and activation functions. Furthermore, the learning rate can be tuned, i.e. the rate every neuron adapts its weight to the overall error to ensure even distribution of learning, as otherwise during backpropagation the last layer, where the error is maximal, would learn the most. To reduce overfitting, large weights can be penalized by a weight decay. Neural networks are very powerful classifiers, but either require expert knowledge or exhaustive search for appropriate hyperparameters to be tuned accordingly.

3.2.5 Support Vector Machines

Support vector machines (SVM) are algorithms for classification and regression not only popular for authentication tasks, but considered as one of the most flexible and effective machine learning tools [39]. Unlike many other approaches, SVMs are not iterative, but solve equations. For classification,

a hyperplane is fitted between the samples to linearly separate the classes from each other, maximizing the distance. SVM is called a large margin classifier. If two classes are not linearly separable, they are projected into higher dimensions, where they possibly can be separated. Because of the implementation details of SVM, the use of the so called “kernel-trick", the computational costly dimensional transformation never has to be executed. This is accountable for the fast training of SVM.

Instead of determining an arbitrary solution, SVMs always find the global optimum. As a classifier that allows for no errors during training and classifies every training sample correctly is very likely to overfit, SVM features a cost parameter C. A small C allows a higher rate of misclassifications during training, which might lead to a simpler, possibly underfit model. Additional tuning parameters are introduced by different kernels, such as radial basis function (RBF), Gaussian, linear, sigmoid, polynomial and other kernels. For optimal fit, parameter grid search with cross validation is recommended.

3.2.6 Generative Model Classifiers

Some authentication approaches use different generative model classifiers (GMC). Instead of modeling the relationship between the observed variables and the target variable (discriminative models), they model the relationship between all variables, i.e. such a model is not only capable of predicting a target variable, but generating new samples. Therefore, GMC could be used to extrapolate data from small datasets. Some examples are Gaussian mixture, hidden Markov and naive Bayes models.

3.2.7 Other Classification Approaches

Sometimes authentication is performed by simply extracting a score, e.g. by cross correlation, similarity or dissimilarity between sample vector and template vector, and comparing it to a threshold. Authenticity is granted, if similarity is above or dissimilarity is below a certain threshold. For identification, the template that results in the highest similarity or lowest dissimilarity score is associated with the sample.

3.3 Others

Some tools and terms are frequently used throughout this thesis, but match non of de above categories.

3.3.1 Tuning Parameters

Classification models typically have one or several tuning or hyperparameters, specifying the learning behavior of the model. Those parameters depend

on problem, data and classification model and need to be specified for every application. Selection of appropriate values for tuning parameters is important for adequate model generalization. If inappropriate values are selected, classification models might over- or underfit training data and therefore won't perform well on real world data.

3.3.2 Parameter Grid Search

One technique of finding appropriate values for tuning parameters is parameter grid search. The goal of parameter grid search is to find tuning parameters, which lead to proper generalization of the model. Classification models are trained and validation results are compared for different values of tuning parameters within the respective typical ranges. For more than one tuning parameter, permutations of different values per parameter need to be evaluated, resulting in n^p permutations, for n being the number of values tried out per parameter, and p being the number of parameters. When typical ranges of certain tuning parameters are unknown, it might be advisable to try powers of ten, e.g. from 10^{-5} to 10^5. Once the order of magnitude is known, optimal parameters can be found by stepwise refining the parameter grid.

3.3.3 Gallery Dependence

The available data for model training is usually limited and sometimes real world data is not available at all. Models are therefore trained on data recorded under different conditions than real world data. This might lead classification models to work well on the training data, but fail to correctly predict classes under real world conditions. Such a model is called gallery dependent. In order to achieve a gallery independent model, training data should contain the same amount of variance than real world data. Furthermore, it is recommended to shuffle the data before data partitioning is done. This way, any bias that was introduced within the recording session is randomly distributed over training, validation and test partitions and therefore affects the model less gravely.

3.3.4 Principal Component Analysis

Principal component analysis (PCA) is a statistical procedure commonly used for dimensionality reduction technique in the machine learning domain [33]. It uses orthogonal transformation, that turns a set of possibly correlated variables into a set of uncorrelated variables of equal length m. This transformation is reversible, as no information has been lost. It can be considered as another representation of the same information. However, the set of variables after transformation has the same length or dimension as the original set. Dimensionality reduction is only introduced, when removing

variables from the set. As we want to reduce dimensions or variables from the set without losing all the information contained, they are not removed arbitrarily. The transformation results in a set of variables ordered according to the variance they represent. The first dimension contains the most variance, while every subsequent dimension contains as much variance as possible, given that it is orthogonal to previous dimensions. Therefore it is possible to represent arbitrary amounts of original variance in $n \leq m$ dimensions.

PCA provides two main benefits. First, dimensionality reduction is very useful, as it facilitates all calculations. Furthermore, the amount of data required for model training l is reduced, as $l >> n$. Second, PCA provides uncorrelated variables, which is beneficial for classification, as correlated variables don't add information to models and increase l.

Chapter 4

ECG as Biometric

In chapter 2, we discussed the capabilities and limitations of state of the art authentication technologies and the need for new authentication methods arising from the unsatisfactory security/usability tradeoff of traditional methods and well known biometrics. In the following chapter, based on the need for secure yet usable biometrics that allow for continuous authentication, we describe ECG as biometric in detail. We start with the properties of ECG waves and proceed with ECG recording methods. Finally, several ECG authentication approaches are presented and evaluated.

4.1 ECG Wave

4.1.1 Characteristics of ECG Waves

An ECG wave depicts the electrical potential of the heart over time [3]. Figure 4.1 shows a schematic of an ECG wave. Each ECG wave represents the activity during one heartbeat. They consist of and can be divided into segments. Each segment of the wave corresponds to one phase of the heartbeat. The amount of heartbeats per minute is commonly known as heart rate or pulse. The duration between consecutive heartbeats is known as RR interval. For healthy people, the meantime between consecutive heartbeats varies even at resting conditions. This is known as heart rate variability (HRV) and depicted in figure 4.2.

P Wave

The heartbeat is initiated by the sinoatrial node, the pacemaker of the heart. It spontaneously generates an electrical impulse that, upon propagation through the heart, causes the cardiac contraction. The P wave typically has a duration of less than 120 ms and a spectrum of 10 Hz to 15 Hz, according to Agrafioti and Hatzinakos [3] and Yanowitz [86].

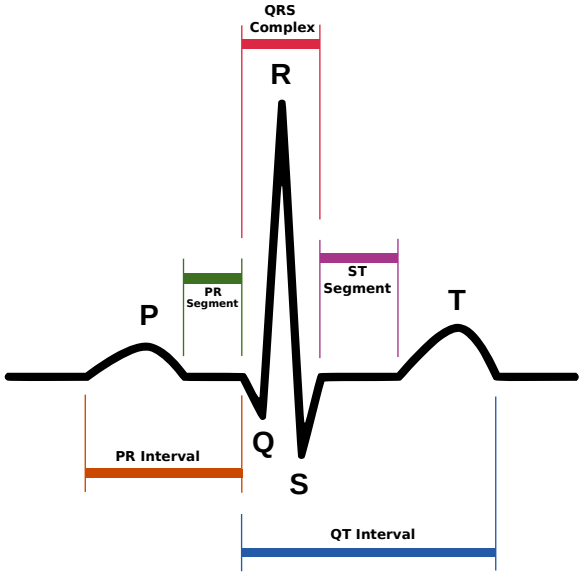

Figure 4.1: Schematic diagram of normal sinus rhythm for a human heart as seen on ECG from [68].

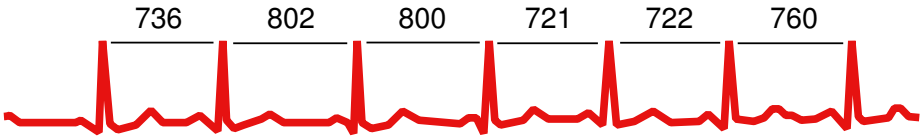

Figure 4.2: Illustration of a healthy ECG from [85]. RR intervals are stated in milliseconds. The heart rate varies.

QRS Complex

After the P wave, usually three consecutive peaks can be observed. The Q, R and S waves correspond to the depolarization of the ventricles, that initiate the contraction of the heart. As depicted in figure 4.1, the downward Q wave precedes the upward R wave, followed by the downward S wave within about 100 ms. The QRS complex has a spectrum of 10 Hz to 40 Hz [3].

T Wave

Finally, repolarization of left and right ventricles correspond to the T wave. It has a duration of about 160 ms and it appears 80 ms to 120 ms after the QRS complex, depending on the heart rate. Agrafioti and Hatzinakos [3] state that, the ST segment is shorter for higher heart rates.

4.1.2 Long Term Changes in ECG Pattern

As the ECG pattern reflects the individual anatomy, it is subjected to a continuous change. Changes in the ECG can be traced back to two main factors. Firstly, the normal process of growing and aging leads to changes visible in the ECG. Secondly, cardiovascular conditions have an immediate effect on the ECG.

Childhood to Adolescence

Carmona et al. [13] observed the development of ECG over a five year period. In 2002, samples of 52 participants aged 19 were taken. In 2007, recordings were repeated on the same group of participants, now aged 24. The study shows a decrease of HRV from early age on. A decrease of HRV can be linked to a loss of complexity of the cardiovascular dynamics [43]. Women and persons who don't practice sports are more affected than men and people who practice sports.

Singh and Gupta [58] conclude, that changes to the ECG due to aging are not consistent. Generally, the heart rate decreases, causing the P wave, QRS complex and PR interval to increase in duration. While the amplitude of the P wave remains consistent over years, the amplitudes of R and S waves decrease.

Advanced Age

The progressive changes in anatomy that affect the ECG mainly happen during childhood and adolescence. Changes of the ECG amongst elderly people often indicate cardiac disorders. Khane, Surdi, and Bhatkar [34] conducted a study with 400 participants aged 45 to 74, all apparently healthy and asymptomatic. 38% of participants showed ECG abnormalities, and prevalence increases with age. The most common abnormalities are left axis deviation (a condition where the mean electrical axis of the heart is misaligned), sinus bradycardia (a lower than normal heart rate), bundle branch block (a disorder of the electrical conduction system of the heart) and ST-T wave abnormalities.

Aging and ECG Authentication

The effects of aging and disease change the ECG considerably. In the ECG authentication toolchain, a change of the ECG pattern would affect the process of feature extraction. Depending on the selected features, aging effects potentially have an negative effect on the classification performance. Different features and feature extractors will be explained and discussed in 4.3. Extractors that rely on the amplitude, position, sequence or interval of ECG

waves are especially prone to abnormalities such as bundle branch block or premature ventricular contraction.

However, the concept of continuous authentication allows for a certain adaptability to long term changes. The repeated authentication reduces the impact of a discrete event caused by cardiac conditions, such as a skipped heartbeat. Furthermore, online learning could be introduced to biometric authentication algorithms. Successful authentications could be added to the positive class, while old samples could be removed from training set after a decay period. This would allow an adaption to long term changes.

4.1.3 Short Term Changes

Multiple factors can influence the ECG, including physical or mental stress, ambient temperature, medical condition, emotional states, etc. In this section, we are going to discuss the effect of physical and mental stress and medical conditions.

Physical Stress

According to Whyte and Sharma [64], physiological stress caused by exercise generates a multitude of responses in the ECG wave. Physical stress leads to an increase in heart rate, with a linear relationship between heart rate and workload. As the duration of the systole (the contraction of the heart) remains constant at about 300 ms, the duration of the diastole (the relaxation of the heart) varies and decreases with increasing heart rate. Further changes of the ECG include changes in amplitude of the P wave, increase in Q wave amplitude, increase in R wave amplitude during medium stress, decrease of R wave amplitude during maximal stress, shortening of QRS complex, changes in T wave and ST segment, decrease of QT interval, as well as superimposition of P and T waves of consecutive beats.

Mental Stress

Taelman et al. [60] studied the effect of mental stress on heart rate and HRV. The autonomic nervous system consists of the sympathetic and parasympathetic nervous system. While sympathetic activity leads to an increase in heart rate, e.g. during sports exercise, parasympathetic activity lowers the heart, e.g. during sleep. The constant interaction of both systems is represented by heart rate variability. When a person is exposed to mental stress, the parasympathetic nervous system is suppressed and the sympathetic nervous system is triggered. Stress hormones such as epinephrine and norepinephrine are released, inducing a 'fight-or-flight' reaction. As a reaction to the hormones, the blood vessels contract, resulting in higher blood pressure. Also muscle tension increases and heart rate rises.

It is concluded that, when exposed to mental stress, subjects show an increased heart rate and a decreased HRV. Although subjects react differently to particular stress situations depending on age, gender and environment, 24 out of 28 subjects show significantly increased heart rates when exposed to mental stress.

When frequency analysis is applied to the signals, the subjects show an increase in the low frequency/high frequency ratio. Low frequency components are associated with sympathetic activity, high frequency components are associated with parasympathetic activity. An increase in the LF/HF ratio indicates higher sympathetic activity, when exposed to mental stress [60].

Medical Condition

As the ECG is a widely used noninvasive medical exploration technology, many medical conditions of the cardiovascular system have an immediate effect on the morphology of the ECG wave. Depending on the kind and severity of the condition and the authentication process, medical conditions could have a negative impact on the authentication performance.

Li and Li [42] mention, that supervised classifiers won't perform well for out-of-set testing samples, and therefore for patients with irregular cardiac conditions. Supervised classifiers would require reenrollment of altered ECG patterns or robust normalization of biased ECG samples. A system, calculating the difference between target and test samples and comparing it to a global threshold is proposed in [42]. Furthermore, beat normalization or outlier removal is needed.

Short Term Changes and ECG Authentication

ECG authentication relies - like all biometric authentication approaches - on the identification, extraction and quantification of domain specific features. Physical or mental stress, as well as environmental factors can influence the ECG wave. While the morphology of the ECG remains comparatively constant, the intervals between subsequent pulses are heavily affected [3]. Heavy stress can lead to changes in amplitude of specific waves, or even superimposition of consecutive beats.

Depending on the feature selection and extraction process, short term changes potentially could have a negative influence on classification performance, specifically on the false negative rate (FNR). Counter measures could include adding ECG samples recorded in stress or post stress situations to the positive class or implementing robust beat normalization and outlier removal to the system. System design heavily depends on the use case and the desired false positive/false negative rate.

4.2 ECG Recording

The first step in an biometric authentication process is data recording. It is required for both training of classifiers and application of authentication systems. Typically, ECG signals are captured by applying sensors to the skin. Recently, sensors emerged that do not require physical contact to the skin. In this section, some of the mostly used sensors are reviewed.

4.2.1 Resistive Electrodes

Resistive electrodes are simple conductive parts that make contact between skin and electrocardiograph. They can be divided into wet and dry models.

Wet Electrodes

Traditional, medical ECG mostly uses silver/silver chloride (Ag/AgCl) electrodes. Those medical standard electrodes need some electrolytic gel between skin and electrode to improve conductivity and therefore are called *wet*. For medical ECG, 10 Ag/AgCl electrodes are placed around the chest and on arms and legs, as shown in figure 4.3. This setup has some severe drawbacks for authentication purposes. Application of electrolytic gel might be perceived inconvenient, as it is time consuming and invasive. Electrode placement around the chest might not be applicable for mobile usage. For chronic, continuous use, the reliability of wet electrodes cannot be guaranteed, as the conductive gel dehydrates over time reducing the conductivity. Smearing of the gel can cause short circuits between adjacent electrodes. Once a recording session is in progress, conductive gel cannot be reapplied without interrupting the record. Furthermore, wet electrodes and conductive gel are reported to have caused irritations on the skin or even dermatitis [55]. Additionally, the signal to noise ratio (SNR) is limited by electrode/skin impedance.

Dry Electrodes

Dry resistive electrodes use the same functional principle, but do not require any conductive gel, fixing the above mentioned drawbacks. They consist of conductive materials such as stainless steel, titanium, aluminum or silver alloys.

The use of dry sensor fundamentally increases flexibility, usability and user acceptance. Yoo et al. [65] include a set of electrodes in a T-shirt by directly printing a planar circuit out of silver-based paste on the fabric. The proposed wearable allows for unobtrusive, long-term ECG monitoring. Matias et al. [46] use a set of dry resistive electrodes placed on the chest, along with capacitive body coupled communications, a wireless body area network to transmit the ECG signal to a portable monitor. Silva et al.

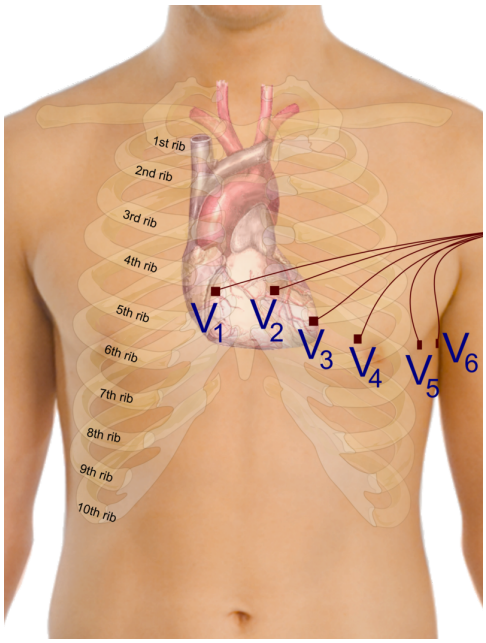

Figure 4.3: Placement of precordial electrodes around the chest from [74]. Additional electrodes are placed on arms and legs.

[57] integrate two dry Ag/AgCl electrodes into a sensor pad, that is placed in front of a PC keyboard and measures ECG between both hand palms while users are typing on the keyboard. Lourenço, Silva, and Fred [44] use three Ag/AgCl electrodes on a sensor mount to capture ECG between left and right hand. The positive and ground electrodes are attached to the left hand index finger, while the negative electrode is connected to the right hand thumb.

By using dry sensors, ECG can be captured nearly anywhere, anytime, unobtrusive and unperceived by others. This greatly facilitates the use and acceptance for both, medical and biometric applications.

However, improved usability comes at the cost of possible drawbacks in signal quality. The electrolytic gel of wet electrodes works as an shock absorbing layer between skin and electrode. The lack of this buffer can cause motion artifacts, i.e. temporary changes of ECG signal caused by motion of users, as skin impedance varies with motion and pressure of the electrode on the skin.

4.2.2 Capacitive Electrodes

Other than resistive electrodes, capacitive or insulating electrodes are not electrically conductive connected to the skin, but rely on capacitive coupling. The surface of capacitive electrodes is covered with a dielectric. Some

capacitive electrodes are even able to capture the ECG through a layer of clothing, making the process of ECG capturing even less obtrusive than dry resistive electrodes. Downsides might be the effect of coupling capacitance on the ECG.

In [50] and [53], three capacitive electrodes are used in wearables. Two sensors are placed on the chest, while one is placed above the pelvis. Special attention is payed in circuit design and component selection, as small size and low power are very important for on-body systems. The proposed system provides signal quality comparable to wet ECG and operates through cotton or wool cloths.

In [14] and [15], capacitive sensors for ECG and electroencephalography (EEG, recording of electrical activity of the brain) are proposed. The systems use an off-body design, allowing for more flexibility in system design and use. The authors report successful ECG recordings through thin shirts as well as through thick sweaters, allowing for maximum comfort during use.

While on-body systems for ECG data acquisition are well documented, McDonald et al. [48] propose a stationary system. Capacitive noncontact sensors are attached to the back of a chair. Subjects only have to sit on a chair and comfortably lean back, without removal of clothing or preparation of skin. While this method implements sensors into infrastructure and therefore doesn't qualify for ubiquitous use, it offers unobtrusive data recording and can be implemented into various environments, e.g. offices or cars.

4.2.3 Comparison of ECG Sensors

The above mentioned examples indicate that capacitive electrodes are better suited for mobile, continuous use than resistive electrodes, especially wet electrodes. In terms of convenience, unobtrusiveness and user acceptance, capacitive electrodes seem to outperform resistive electrodes. However, Searle and Kirkup [55] conducted a comparative study on the performance of ECG sensors. In this section, properties of wet, dry and capacitive sensors are compared based on the findings in [55].

Power Line Noise

One major interference to bioelectrical signals is power line noise, a 50 Hz to 60 Hz noise induced by the omnipresent power supply. Power line noise affects ECG electrodes in two ways. First, capacitive coupling induces voltage on the leads of the circuit. Generally, nearby leads are similarly affected by power line noise, so no voltage difference can be measured between them. But different impedance in two leads, e.g. caused by different skin/electrode impedance leads to a differential voltage interfering with the ECG. All electrodes are affected by this effect. Second, power line noise is induced in the human body as a common mode signal. Different contact impedance

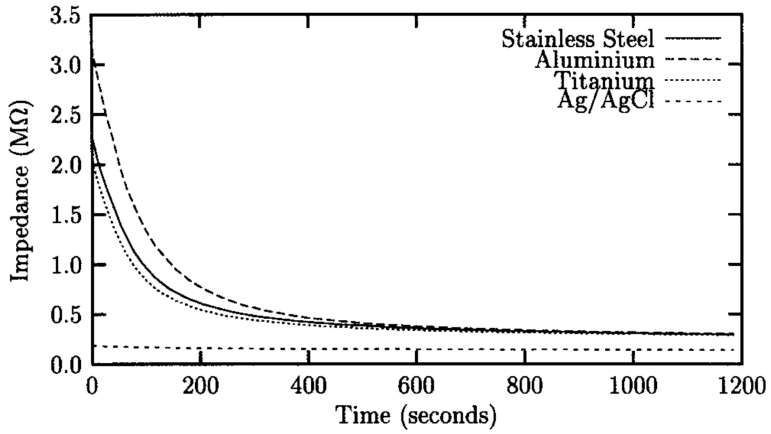

Figure 4.4: Findings of Searle and Kirkup [55] regarding the contact impedance of different dry electrodes, compared to standard Ag/AgCl wet electrodes.

leads to a voltage divider effect, causing differential voltage at the electrodes and therefore interfering the ECG. Counter measures include reducing the output impedance of the electrodes, as for low impedance the difference in electrode impedance and therefore the induced voltage difference is also low. This can be achieved by using active electrodes, buffered by an operational amplifier with unity gain. Active electrodes still suffer from power line noise, due to the capacitive coupling on the circuitry before the buffer and finite impedance of the operational amplifier, but about two orders of magnitude lower than conventional, passive electrodes.

Impedance over Time

As capacitive electrodes have a constant impedance, mainly depending on the operational amplifier used, resistive electrodes are subject to constant change in impedance. Figure 4.4 shows the decrease of contact impedance over time. The impedance appears to decrease exponentially to a comparative level for all tested dry electrodes. This effect is explained by perspiration that reduces the impedance of the skin. Wet electrodes have a significantly lower impedance, but for long term use, the electrolytic gel dehydrates and contact impedance increases.

Effect of Electrical Charges

As mentioned earlier, electrically charged bodies near the electrodes have an influence on the measurement. Capacitive sensors are expected to suffer more from effects by electrically charged bodies than resistive sensors. Searle

and Kirkup [55] placed a rotating metal rotor, charged with 4 kV over the sensors mounted on the subjects arm. Interestingly, wet electrodes without shielding suffered the most interference. Unshielded, capacitive electrodes do slightly better, and dry resistive sensors suffer the least. When the electrodes are shielded by a grounded, metal surrounding, the interference is reduced by 26 dB on average.

Movement of Electrodes

Moving electrodes have a negative influence on both, resistive and capacitive electrodes. As wet resistive electrodes often are fixed on the subject by an adhesive and due to the electrolytic gel, ohmic contact is maintained between electrode and skin even when mild force is applied to the electrodes. Nevertheless, pushing, pulling or stretching of the skin causes the skin potential artifact, an voltage between inner and outer layers of the skin. For capacitive electrodes, moving the electrodes corresponds to changing the thickness of the dielectric between the plates of a capacitor. Wet electrodes are least affected by motion effects, dry resistive electrodes suffer the most.

Conclusion and Choice of Sensor

We conclude that for our desired use case of mobile, unobtrusive, continuous ECG recording, capacitive electrodes are an appropriate choice. While we cannot identify any grave disadvantages of capacitive electrodes, they offer high tolerance against electrode movement. Furthermore, capacitive electrodes seem to be most usable within this comparison, as the skin doesn't need to be prepared by applying conductive gel or shaving for ohmic contact and can operate through layers of clothing.

4.3 ECG Authentication

In this section, we investigate some of the available techniques for ECG authentication, from feature selection to extraction and classification. Generally, ECG authentication techniques can be divided in two classes: algorithms based on fiducial features and algorithms based on nonfiducial features. Although the name suggests otherwise, many algorithms based on nonfiducial features also rely on the detection and extraction of fiducial features, but they use features other than fiducial points of the ECG for authentication.

4.3.1 Algorithms Based on Fiducial Features

Many authentication techniques use fiducial features of the ECG to distinguish individuals. Fiducial features are based on fiducial points, such as

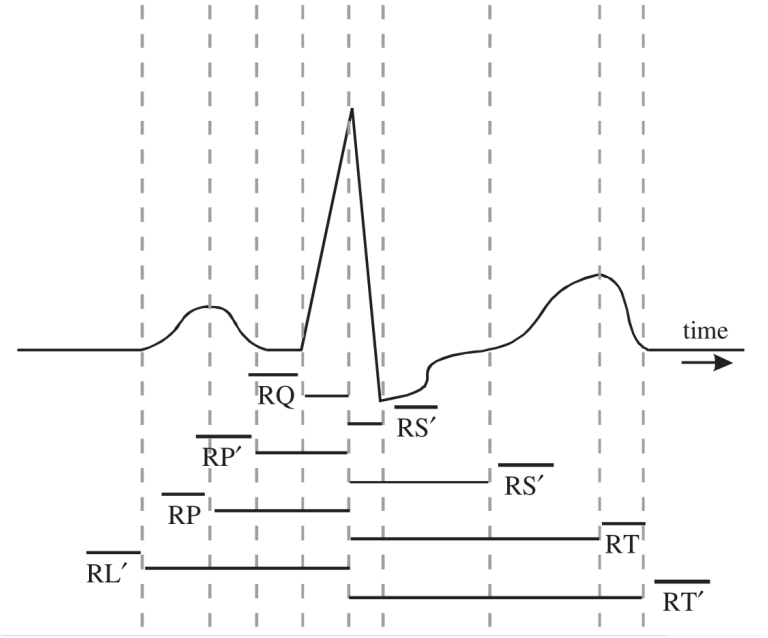

Figure 4.5: Fiducial features from [29]

positive or negative peaks of ECG waves and the relation between them. Odinaka et al. [51] divide fiducial features into five types: temporal, amplitude, area, angle and dynamic (across heartbeats, such as RR interval).

As an example, we have a closer look at the work of Israel et al. [29]. After data recording and filtering, the ECG signal is subjected to fiducial detection. The P, R and T peaks are easily detected, as they are local maxima of the ECG. The onset and offset of the waves are determined by the minimum radius of curvature, as this method proved more robust to noise than using the derivative. Figure 4.5 shows the fiducial features extracted in [29]. As depicted, the authors rather use time intervals between on- and offset and peaks of waves than amplitudes as features. They explain their choice by the invariance of temporal features against sensor position. Unlike amplitude features, they are not affected by sensor placement. Unfortunately, temporal features suffer from changes in heart rate and HRV, as discussed in section 4.1.3. Hence, normalization of the ECG according to heart rate is required. Due to the physiology of the heart, not all segments of the ECG are equally affected by a change in heart rate, see section 4.1.1. Therefore, only the segments mainly responsible for heart rate changes, namely P and T complexes are normalized by dividing by the heart rate. The current heart rate can be determined by the L'T' distance. Classification is performed by linear discriminant analysis (LDA) on the acquired features. Israel et al. [29] claim to reach $97 - 98\%$ classification performance upon their dataset.

However, fiducial-based approaches rely on exact and robust detection of fiducial points or wave segments. Even a slight error in wave detection or alignment can lead to misclassification. Fiducial detectors usually are built and rely on a healthy ECG without abnormal findings. It is unclear, how fiducial detection algorithms would perform on patients with irregular cardiac conditions, such as premature ventricular contraction or under the influence of heavy stress, that causes superimposition of consecutive beats.

4.3.2 Algorithms Based on Nonfiducial Features

Algorithms based on nonfiducial features don't use fiducial points for feature vectors, but usually split the ECG sample into overlapping or nonoverlapping windows and extract features from those windows. As an example, we have a closer look on the method proposed by Agrafioti and Hatzinakos [3].

The first step is the application of a bandpass filter to remove low-frequency noise, such as baseline wandering and high-frequency noise, such as power line interference. Then, the filtered ECG signal is cut into nonoverlapping windows. The only constraint during windowing regards the window size. Each window must at least contain one full heartbeat; in [3] a window size of 5 seconds is used.

Now, autocorrelation is applied on the windows with $M << N$, where M is the time lag and N is the window size. Autocorrelation provides the main advantage of this method, as the resulting signal is already normalized to the maximum correlation at time shift 0.

For dimensionality reduction, DCT or LDA is applied. From the resulting signal of length $N+M$, only $C << M$ non zero DCT or LDA coefficients that contain the most significant information are selected. Validation is done by a threshold on the euclidean distance between the sample and the positive class, therefore the class which is meant to be evaluated positively by authentication and contains samples of the legitimate user. If the distance exceeds a certain value, authentication is denied. The false positive rate and false negative rate depend on and can be influenced by the threshold.

Agrafioti and Hatzinakos [3] report a classification rate of $96 - 100\%$ on their dataset. The simple yet robust design of the algorithm seems to match the requirements of continuous authentication, as it doesn't depend on fiducial feature extraction. Further, the validation by euclidean distance and threshold could be exchanged and possibly improved by machine learning classifiers.

4.4 Evaluation of Related Work

We have discussed the properties of ECG waves (section 4.1), how they can be captured (section 4.2) and how recognition is performed on this data (section 4.3). In the following section, we summarize and discuss the

different approaches, their advantages and disadvantages as well as their discriminative power in terms of authentication.

4.4.1 Data Acquisition

If ECG data should be used for authentication, the first step is record data. Currently, there are three types of electrodes available to capture ECG on the skin. Medical standard Ag/AgCl electrodes require electrolytic gel to reduce the sensor/skin impedance. The gel also acts as a mechanical buffer against vibrations and the usually self-adhesive electrodes are comparatively invariant against movement. Good signal quality comes at the price of inconvenient use. Preparations include undressing of the subject, application of conductive gel, possibly shaving the contact points on the skin and application of the electrodes. Wet electrodes are hardly usable for authentication means.

Dry resistive electrodes do not use conductive gel, but as they too rely on ohmic contact to the skin, possibly require preparations such as undressing and shaving. Usability and flexibility are greatly improved compared to wet electrodes, as they can be implemented into wearables [65] or ECG can be captured unobtrusively from the subjects hands [44, 57].

Capacitive or insulating electrodes don't rely on ohmic contact, but on capacitive coupling. This allows for ECG recording even through layers of clothing [14, 15] and without further preparations. While signal quality and invariance against the movement of wet electrodes is unreached, capacitive electrodes clearly outperform resistive electrodes in terms of usability. Unobtrusive data recording is crucial for usability and user acceptance, therefore we consider capacitive electrodes as key technology for ECG biometrics.

4.4.2 ECG Features

Once the data is recorded, a set of features needs to be extracted, in order to perform authentication. We can distinguish fiducial and nonfiducial features. While fiducial features are derived from characteristic points in the ECG wave, e.g. peaks or sections of waves, nonfiducial features are extracted from ECG signals after further processing such as segmentation or cross-correlation. To extract features from fiducial points, those points need to be detected first. This can be a challenging task, as a standardized definition of fiducial points in the ECG wave is yet to be found. Furthermore, conditions like electrode placement have an inherent influence on the signal amplitude and therefore on amplitude features. Additionally, temporal changes in the ECG wave, like an increased heart rate, or medical conditions can severely influence temporal features. Nonfiducial features, e.g. derived from autocorrelated windows are possibly less affected of temporary changes in amplitude or time.

4.4.3 ECG Classification

After recording the ECG signal and extracting the selected features, classification is performed. In the simplest case, a distance metric is employed to calculate the difference or similarity between sample and reference vector. Class membership is assigned based on whether the similarity or dissimilarity is above or below a certain threshold. Other classification methods frequently used in literature include KNN, LDA, neural networks, SVM, etc.

4.4.4 ECG Authentication in Literature

Now that we summarized the building blocks of ECG authentication systems, we discuss several approaches proposed in literature, that serve as the foundation of our own work.

Shen, Tompkins, and Hu [56] describe a system that extracts seven fiducial features from the preprocessed ECG signal. The selected features include temporal and amplitude features from the QRS complex, as well as the QT interval normalized with the heart rate. The features are subjected to a two-stage authentication process. First, a template matching algorithm is applied. For every subject, a template waveform is selected from the signal. During testing, cross correlation is applied on the template and 20 heartbeats of the sample. If the average of the correlation coefficients exceeds a value of 0.85, the algorithm proceeds with stage two. Otherwise, the algorithm concludes that the sample doesn't match the template. A decision-based neural network (DBNN), trained with 20 heartbeats per class is employed to classify the sample. On their database, consisting of 20 subjects, an identification rate of 95 % for template matching and 80 % for DBNN is reported, when applied separately. When DBNN is executed after template matching, the authors claim 100 % authentication rate on the dataset.

Another fiducial-based approach was developed by Israel et al. [29]. Data recording of 29 individuals was split into 12 repeat sessions with seven tasks per session. The tasks consist of meditative and recovery tasks, reading aloud, mathematical manipulation [sic] and driving in virtual reality and were designed to stimulate different states of anxiety. Additionally, data was recorded from neck and chest. After preprocessing, 15 temporal fiducial features were extracted, but only 12 features were selected for classification, that has been done using LDA. The results show that, because of the selection of temporal features, all subjects could be identified, invariant of the sensor location. The classification performance within anxiety state is specified with 97 %, and 98 % between anxiety state.

A partially fiducial approach is adopted by Silva et al. [57]. For preprocessingc the data is filtered by a bandpass filter to remove unwanted noise. Then, segmentation based on R-peak detection is used for feature extraction. Additionally, outlier removal is performed by calculating the mean cosine

distance to the average heartbeat waveform of the sample. The database includes ECG recordings from 63 individuals, recorded in two sessions 4 months apart. In their experiments, two different classification approaches are tested for both, within and between session authentication. A KNN classifier for cosine and Euclidean distance metrics is employed, as well as a SVM classifier. The ECG waveform template is acquired either by the mean or median heartbeat waveform. Within all tested permutations, SVM classification serves the best results. For within session authentication, an equal error rate (EER) of 0.99 % for the first recording session, and 1.92 % for the second recording session was achieved. Between session authentication reached an EER of 9.1 % and 9.37 %, when one session is used for training and the other for testing.

A similar approach was presented by Lourenço, Silva, and Fred [44]. As usual, the signal is first filtered to remove non signal components. Detection of QRS complex is used for segmentation. Other than the approaches presented so far, the fiducial features used for classification are the amplitudes of the mean waveform. The direct use of the waveform makes normalization necessary, therefore both, time and amplitude normalization are applied on the signal. For classification, KNN with $K = 1$ and Euclidean distance criterion is used. Within a database of 16 subjects, 94.3 % identification rate and for authentication, an EER of 10.1 % was achieved.

Another promising authentication technique is proposed by Fatemian and Hatzinakos [18]. Preprocessing is performed after discrete wavelet transformation and reconstruction using the quadractic spline wavelet. After analyzing the wavelet, the authors concluded that most energy of the ECG waveform is contained in the 3^{rd} scale, while removing effects of high frequency noise and power line interference. QRS complex, P and T wave are detected and resampled for heart rate, zero-mean and unit variance normalization. For each class, a template is obtained by calculating the median of the normalized heartbeats. For classification, an iterative template matching algorithm is employed. The number of heartbeats used for the extraction of an ECG template is increased every iteration, if a correlation threshold between sample and reference ECG template cannot be met. If after 10 iterations the similarity threshold hasn't been exceeded, the system fails to recognize the subject. A high threshold corresponds to a low false positive rate while a low threshold corresponds to a low false negative rate. Within 27 subjects, an identification performance of 99.63% was achieved.

Regarding the high effort necessary for peak detection, segmentation and normalization, a much easier approach to the problem is presented in [3]. After filtering the data, the signal is cut into nonoverlapping windows. Segmentation is performed without prior peak detection and contains at least one heartbeat. The normalized autocorrelation of each window is calculated. Then, DCT or LDA are applied for dimensionality reduction and template matching is done using Euclidean distance. The experiments were performed

Pub.	n	Feature Type	Feat	l	WIR	SIR	EER
[56]	20	fiducial	7	20 HB	-	100%	-
[29]	29	fiducial	12	120 sec	63-83%	97-98%	-
[57]	63	part. fiducial	-	5 HB	-	-	0.99% - 9.37%
[44]	16	fiducial	-	30 HB	-	94.3%	10.1%
[18]	27	fiducial	-	10 HB	-	99.63%	-
[3]	27	nonfiducial	-	5 sec	86.3-95.9%	96.3-100%	-

Table 4.1: Comparative table of ECG authentification approaches in literature. n=number of individuals in database, l=sample length, Feat=number of features, WIR=Window Identification Rate, SIR=Subject Identification Rate, HB=Heartbeats.

on the same dataset of 27 subjects as in [18]. For DCT, 96.3% and for LDA 100% of subjects were identified correctly.

The above mentioned algorithms show promising results on their respective datasets. Nevertheless it is hard to compare them, as ECG recording conditions as well as database size vary. Table 4.1 shows a summary of the main characteristics of the mentioned authentication approaches. Note that in [57], the EER is mentioned instead of the overall identification performance. In [3, 18, 44, 57], no explicit feature extraction is performed. The ECG wave or derivatives or transformations of the ECG wave are used as features or an unspecified number of features is selected. The window identification rate stated for [29] and [3] is of importance, as it represents the identification performance under a time constraint. If sufficient data and time is available, even a WIR of slightly over 50% could correctly classify a sample by majority voting and therefore achieve a high SIR. For approaches that use discrete time frames in the dimension of one or several heartbeats, especially for continuous authentication, this real time classification property would be of great importance and probably more important than the overall performance.

4.5 Potential of ECG for Continuous Authentication

Summarizing, ECG depicts the sum of the electrical potential of the heart over time, differences in individual physiology become visible in the characteristic waveform measured on the skin. Features used for recognition include fiducial points, such as on- and offset of waves, peaks or features derived from those points, like amplitudes or intervals, as well as statistical features. Data acquisition setups are available in a variety of configurations. While medical ECG requires 10 wires placed on chest, arms and legs, which results in a level of intrusiveness and usability even far worse than retina recognition, many ECG systems use wearable or integrated acquisition systems that allow for

unobtrusive, continuous recording. When built into a shirt, watch, bracelet, fitness tracker, smartphone or any other mobile device, usability can even be on par with e.g. gait recognition, as user interaction is limited to carrying the sensor device. Security-wise, ECG recognition systems are reported to provide considerable recognition rates [51], given a well suited recording setup. Furthermore, unlike many behavioral biometrics, ECG can't be easily mimicked by attackers and recording requires close distance to the subject, which makes it harder to be disclosed, compared to obvious and publicly visible biometrics like face or iris.

Based on the findings presented in section 4.4.4, we estimate the biometric potential of ECG similar to face recognition, while potentially providing a high level of usability and allowing for continuous authentication. ECG provides advantageous properties for authentication and together with other biometrics, it can provide strong authentication.

Chapter 5

Continuous ECG Authentication System Design

In section 2, we discussed the security issues of session based authentication approaches and the inconvenience associated with continuously entering knowledge-based secrets. Continuous, biometric authentication is one way to combine security and usability for mobile authentication. Among the multitude of available biometrics, ECG has special properties, as stated in chapter 4. In this chapter, we present our prototypic system to continuously capture ECG signals and perform user authentication. As shown in figure 5.1, our approach is divided into five components. After data recording, digital filters are applied to remove unwanted noise from the ECG signal. During segmentation, the signal is cut into windows. Afterwards, a feature extraction stage includes autocorrelation of the ECG windows and calculation of difference feature vectors. Finally, classification models are employed for authentication and identification.

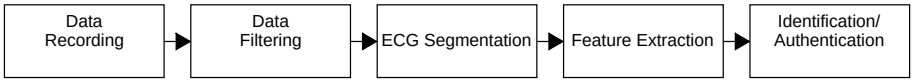

Figure 5.1: Main components of our continuous ECG authentication system.

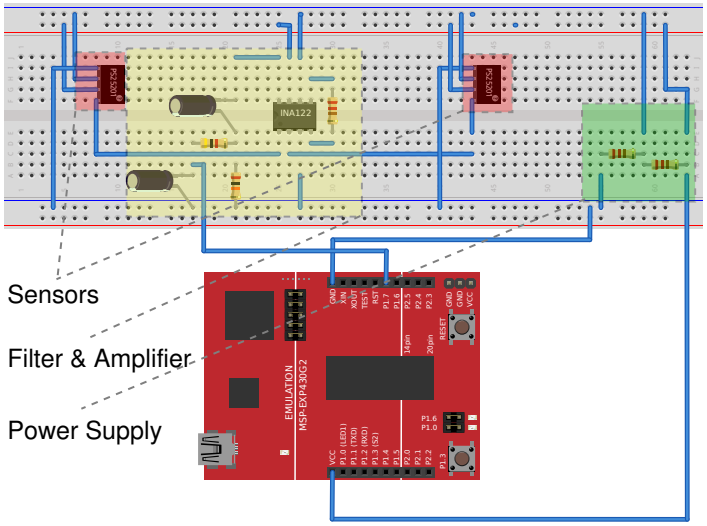

Figure 5.2: Sketch of hardware assembly, created with Fritzing [36].

5.1 Data Recording

In order to perform biometric authentication, ECG data is captured, recorded and transmitted to a processing unit, which performs all subsequent tasks. Our system therefore consists of ECG sensors, which capture the electrical signal on the skin, a microcontoller, that samples the continuous signal and a WIFI interface, to transmit the data to a processing unit. In the following section, our data recording system is explained in detail.

5.1.1 Hardware

Figure 5.2 shows a sketch of our hardware assembly with the sensor, filter and amplifier and power supply regions highlighted. The according schematic is provided in figure 5.3. The components were selected and the system was designed to be mobile, versatile and usable. Use cases include the implementation into a smartphone case, wearable, smartwatch or fitness tracker, etc. Figures 5.4a and 5.4b depict two possible applications of our system.

Data Capturing

Our data capturing system consists of four components. Each of those is described in this section.

Sensor Based on the findings in section 4.2 and the desired use case of mobile, continuous authentication, we chose to use active, capacitive sensors. PS25201 EPIC Ultra High Impedance Electrophysiological Sensor [83] from

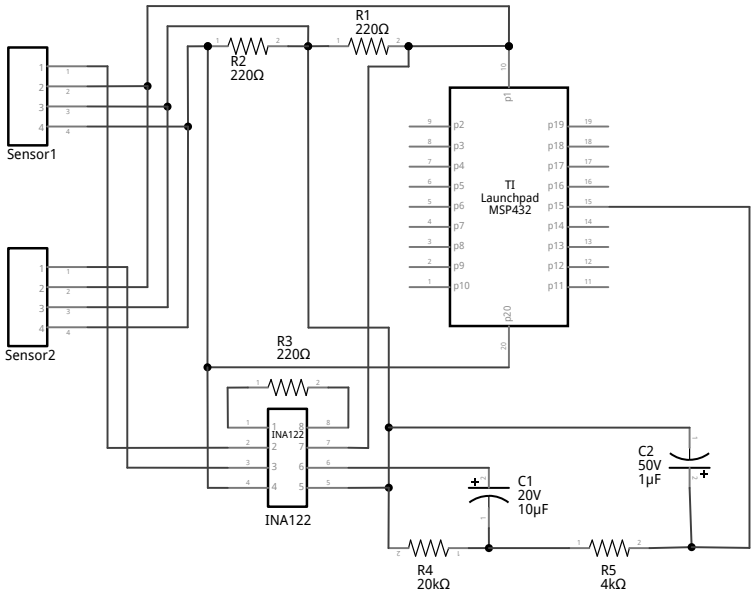

Figure 5.3: Schematic of hardware assembly, created with Fritzing [36].

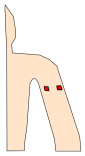

(a) A pair of sensors implemented in a shirt or elastic band is placed on the arm, as explained in [82].

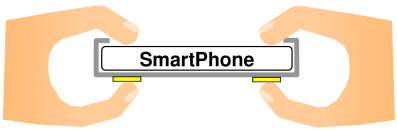

(b) A pair of sensors implemented into a smartphone case, as reported in [81].

Figure 5.4: Use cases for capacitive ECG sensors.

Plessey Semiconductors Ltd. met our requirements. It is hereinafter referred to as *the sensor*. The sensors have a length and width of about 1 cm, and a height of about 3.5 mm, which makes them conceptually capable of being integrated into most mobile devices. They can be operated with a bipolar power of ±2.4 V to ±5.5 V, which allows for a variety of power supplies. The typical power consumption is 2.5 mAh per sensor, which results in a total power consumption of 0.025 W at 5 V. This is sufficiently low to provide convenient battery life for mobile use cases. However, the sensors offer a voltage gain of 50. As the amplitude of ECG signals on the skin is about 1 mV, the sensors have output signal amplitudes of about 50 mV.

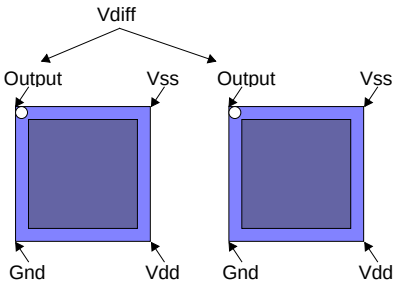

Figure 5.5: The ECG signal is acquired by differential voltage between two PS25201 sensors, adapted from [83].

Amplifier In order to reach the full scale range (FSR) of the analog-to-digital converter (ADC) of 3 V, in the next step, the ECG signal needs to be amplified. Note that FSR is not limited by the microcontroller, but by firmware created with Energia framework [73], which only allows for reference voltage of 3 V. However, is without relevance, as we are also limited in amplification by voltage supply of ±2.5 V. As figure 5.5 shows, the ECG signal is acquired by measuring the differential voltage between our sensors. Therefore, we use an instrumentation amplifier INA122 [75] and set the gain to a factor of 105 by connecting the R_G pins with a 2 kΩ resistor. During system design and testing, this was largest possible gain without suffering from saturation and therefore information loss of the signal.

Filter The now amplified signal suffers from both high and low frequency noise, as described in section 4.3. Therefore, a simple RC bandpass is employed to remove low frequency noise such as baseline wandering and high frequency noise like power line interference. The lowpass filter consists of a 4 kΩ resistor and a 1 µF capacitor, resulting in a cutoff frequency of 40 Hz. The highpass filter consists of a 20 kΩ resistor and a 10 µF capacitor, resulting in a cutoff frequency of 0.8 Hz. The filters have been designed in order to remove unwanted noise, while keeping sufficient information for classification of ECG signals.

Microcontroller

For sampling of the continuous ECG signal captured by the sensors, we use a MSP-EXP432P401R LaunchPad [76] from Texas Instruments. It features a 48 MHz low power CPU and 14-bit ADC, which should suffice for WIFI data transmission at about 100 Hz. The launchpad application board allows for quick and easy prototyping and the extension with further hardware.

WIFI Data Transmission

To transmit the sampled ECG signal to a mobile device, we use a CC3100 wireless plug-in module [77] for the MSP-EXP432P401R LaunchPad. The WIFI booster pack is simply attached to the Launchpad by plugging it onto the pins.

Ground

When the output of our continuous ECG recording system is monitored with an oscilloscope, a 50 Hz rectangular signal can be observed. We traced the rectangular signal back to the capacitive sensors, which have a fixed voltage gain of 50. It is a result of capacitance coupling of circuitry and human body with surrounding alternating current (AC) system, devices and wires, which are present in most buildings, especially our laboratory and office environment. The reason for the waveform to be – other than expected – rectangular instead of sinusoidal, is saturation of the amplifier on the sensors. The high amplitude of coupled noise causes the amplifier to saturate and cut off higher amplitude signal components. As amplifier saturation means information loss for our signal, we need to bypass this effect. One countermeasure is to shield our system against coupling capacitance. As our experimentation setup is built upon a breadboard and a pair of PS25101, a shielded and wired variant of the PS25201 retails at about 1000 € [72], shielding would be a difficult and expensive task. In search of a more practical way to circumvent this effect, we added a connection between subject and ground. Like all conductive connections, it relies on a low impedance interface, as explained in section 4.2. For availability and cost efficiency, we asked the participants to keep a spoon connected to ground in their mouth. This additional wet electrode successfully eliminated the effect of coupling capacitance from the AC system.

However, an additional ground electrode drastically reduces usability and practicability of an continuous, mobile ECG authentication system. Nonetheless, this doesn't reduce the potential of continuous, mobile ECG authentication in any way, as a system design aware of this problem could maintain a two-electrode-setup by shielding sensors and system. By using properly designed hardware, ECG can be unobtrusively recorded and keep a low profile, ideally even unnoticed by the user.

5.1.2 Software

The firmware for our ECG hardware has been created using Energia [73], Release 0101E0017, an open-source platform for electronics prototyping based on Arduino IDE [67]. Energia supports the MSP432 board as well as the WIFI boosterpack CC3100 with pin mappings readily available. Our code consists of a *setup* method, which is executed once to initialize the system,

and a *loop* method, which is executed continuously.

Listing 5.1: Initialization of variables and specification of network SSID and password, as well as host IP and port.

```
// Import Libraries
#include <WiFi.h>
#include <WiFiClient.h>

// variable initialization
int ecg = 0;
char ssid[] = "WIFI_SSID";
char password[] = "WIFI_PASSWORD";
IPAddress server = IPAddress(192, 168, 1, 101);
int port = 9999;
WiFiClient client;
```

Listing 5.2: Firmware setup method. System connects to WIFI first, then a connection to the host is established.

```
void setup()
{
  Serial.begin(115200);
  delay(1000);
  Serial.println("Connect to WiFi network...");
  WiFi.begin(ssid, password);

  while(WiFi.status() != WL_CONNECTED)
  {
    delay(100);
  }

  Serial.println("Connected");
  Serial.println("Acquire IP address...");

  while(WiFi.localIP() == INADDR_NONE)
  {
    delay(100);
  }

  Serial.println(WiFi.localIP());
  Serial.println("Connect to server...");
  client.connect(server, port);
  Serial.println("Connected");
}
```

At first, variables are initialized as shown in listing 5.1. WIFI-SSID, password, host address and port need to be specified to establish a socket connection. After variable initialization, the *setup* method depicted in listing 5.2 is executed. The system tries to connect to the WIFI network and waits until the connection is established. As soon as the DHCP sever assigned an

Listing 5.3: Firmware loop method. The system repeatedly reads the input voltage from our ECG sensors and transmits it to the host.

```
void loop()
{
  delay(10);
  ecg = analogRead(30);
  client.println(ecg);

  if(!client.connected())
  {
    client.stop();
    client.connect(server, port);
  }
}
```

IP address to our system, it connects to the host specified above by IP and port. The *loop* function contains the duty of our system (listing 5.3). The ADC reads the current analog voltage from the input pin connected to our sensors and assigns the corresponding digital value to the variable *ecg*. In case the connection to the host is interrupted, the system tries to reconnect. Data acquired during connection downtime is not buffered, which causes missing data when connection is lost. However, although a system ready for serial production might would contain a buffer, for real time authentication it is not necessary. When connection to the sensor is lost, authenticity of users cannot be verified. Therefore, negative authentication attempts as result of missing data can be considered as intended behavior. Every cycle, the loop is delayed for 10 ms to achieve a sample rate of about 100 Hz. More preciscely, for a desired frequency band from 1 Hz to 40 Hz, a sampling rate above 80 Hz is required, according to the Nyquist-Shannon sampling theorem [32]. As the number of statements within the loop is reasonably low, their execution time can be neglected when calculating the sample rate. We measured the sample rate by calculating the interval time between consecutive messages from the system, resulting in an average sample rate of 95 Hz. We conclude an execution time of 10.52 ms for one loop cycle. By selecting a delay of 9 ms, the execution time would have been 9.52 ms, leading to a sample rate of 105 Hz. As 95 Hz is still sufficiently far above the minimal sample rate of 80 Hz, we decided to stick with a delay of 10 ms.

5.2 Preprocessing

In the next step, the samples recorded by our system and aggregated by a desktop client are digitally filtered, cut and selected. Digital filtering and manual selection are performed on a desktop computer. Digital filters have been created and applied using Matlab's filter designer [78]. For highpass

filtering, we designed a Butterworth filter with a stopband frequency of 0.1 Hz and a passband frequency of 1 Hz, as shown in listing 5.4. For lowpass filtering, we designed a Butterworth filter with a passband frequency of 5 Hz and a stopband frequency of 10 Hz, as shown in listing 5.5. Although the cutoff frequency seems to be very low, this filter effectively removes high frequency noise while recovering a distinctive ECG waveform.

Listing 5.4: Highpass filter designed with Matlab's filter designer [78].

```
Fs = 100;  % Sampling Frequency

Fstop = 0.1;          % Stopband Frequency
Fpass = 1;            % Passband Frequency
Astop = 80;           % Stopband Attenuation (dB)
Apass = 1;            % Passband Ripple (dB)
match = 'passband';   % Band to match exactly

% Construct an FDESIGN object and call its BUTTER method.
h  = fdesign.highpass(Fstop, Fpass, Astop, Apass, Fs);
Hd = design(h, 'butter', 'MatchExactly', match);
```

Listing 5.5: Lowpass filter designed with Matlab's filter designer [78].

```
Fs = 100;  % Sampling Frequency

Fpass = 5;            % Passband Frequency
Fstop = 10;           % Stopband Frequency
Apass = 1;            % Passband Ripple (dB)
Astop = 80;           % Stopband Attenuation (dB)
match = 'passband';   % Band to match exactly

% Construct an FDESIGN object and call its BUTTER method.
h  = fdesign.lowpass(Fpass, Fstop, Apass, Astop, Fs);
Hd = design(h, 'butter', 'MatchExactly', match);
```

Figure 5.6 shows the acquired signal of several heartbeats as transmitted by our system in contrast to the same sample after digital bandpass filtering. As visible on the y-axis, filtering introduces shift and scale to the signal. The amplitude of the filtered signal is lower than the original amplitude, and the signal is centered around 0 V, as highpass filtering removes any constant signal components, i.e. signal parts with 0 Hz or voltage offsets.

In the next step, signals undergo a manual selection and cutting process. As described in chapter 6, ECG recordings were sometimes interrupted or influenced by movement of the participants or abrupt changes in the impedance of the interface between participant and ground. Therefore, each recording was monitored and heavily biased sections were manually removed. Additionally, the recordings are cut to equal length. We do this to maintain

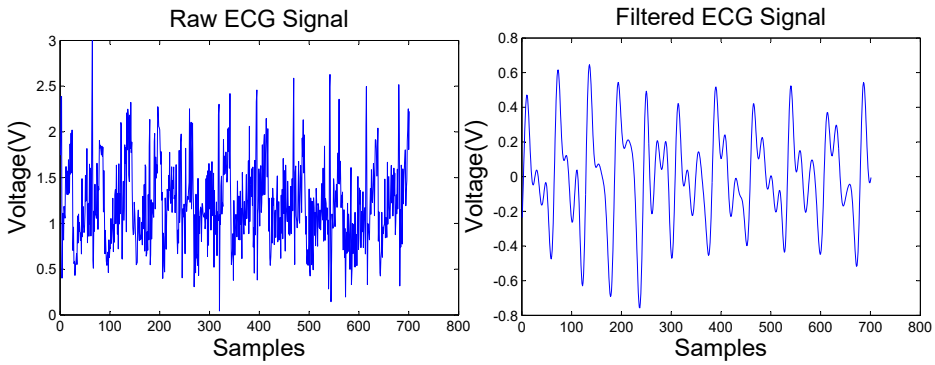

Figure 5.6: The left chart depicts the sampled ECG signal, as transmitted by our system. The right chart shows the same sample after digital filtering.

class balance during evaluation. If certain classes are over- or underrepresented in the database, classification results would be biased. The length of the ECG recordings therefore is determined the shortest one. For our evaluation, recordings are cut to a 11000 samples, which matches a duration of 110 s at 100 Hz sample rate.

While manual selection and cutting of ECG signals are conducted for evaluation of our system only and therefore would not be included in productive versions, digital filtering would be performed on mobile devices.

5.3 Feature Extraction

For feature extraction, we employed a variant of the method proposed in [3] that we discussed in detail in section 4.3.2. The featureless approach segments ECG signals into nonoverlapping windows and performs autocorrelation on each of those windows. It has several advantages over other feature extractors, as it doesn't depend on signal alignment, temporal or amplitudinal normalization. Windows can be extracted arbitrarily from a time series, as long as they contain at least one full heartbeat. Using nonoverlapping windows of length n in seconds, a new window is available every n seconds. In case of an attack, e.g. robbery of a currently used and unlocked mobile device, it would take on average $n/2$ seconds until the next authentication and after n seconds, the next window doesn't contain any ECG data of the user anymore.

For continuous ECG authentication, a sliding window can be employed, and authentication frequency can be tuned according to the security and real time demand of user and application. For the windows of length N, autocorrelation with $M << N$ timelags is conducted. As a result of autocorrelation, the calculated samples of length $M+1$ are normalized regarding

time and amplitude. To further reduce dimensionality, we employ PCA after autocorrelation. Furthermore, PCA provides us with uncorrelated features, which is beneficial for classification.

5.4 Classification

For classification of the acquired features we use machine learning classifiers. Classification consist of two stages, namely training and testing. During training, a portion of data called training partition is passed to the classifier. We fit classification models to the training data which learn the relations between features and target variable. If tuned correctly, they can predict class membership based on yet unseen real world data. To estimate the performance of the achieved models, we use their predictions on our test partition. The predicted classes are compared to the observed, real classes. Based on accordance respectively difference between predicted and observed classes, performance measures such as accuracy or Cohen's Kappa [25] are used. Kappa is a more robust performance metric compared to simple accuracy, as it incorporates correct classifications by chance. It is defined as

$$\kappa = \frac{p_o - p_e}{1 - p_e}$$

where p_o is the accordance between prediction and observation and p_e is the probability of an accordance by chance. In contrast, accuracy is defined as

$$Accuracy = \frac{p_o}{n}$$

where n is the number of predictions and observations. As an example, we imagine a test for a certain disease. From 100 patients, only 1 person is is sick, while the remaining 99 patients are healthy. If the test would predict all 100 patients to be healthy, it would be right in 99 cases. Accuracy for this example therefore is 99 %. In contrast, Cohen's Kappa evaluates to 0. Therefore, we always mention accuracy combined with Kappa during performance evaluation.
Based on those metrics, we evaluate classification performance of our models. We distinguish between classification for identification and authentication purposes.

5.4.1 Identification

During identification, identity of subjects is predicted based on ECG samples. For each sample, the class is determined which is most likely to match. Identification is a comparison of a sample against every known class within a database. It can only correctly predict classes present in the training set. This requires the training set to contain data of all subjects present in the

test set. Test samples outside the training set would be falsely assigned to an existing class within the training set, causing misclassifications. Therefore, if new subjects should be identified, reenrollment of training data and retraining of the model is required.

5.4.2 Authentication

For authentication, the process of identification is slightly altered. Authentication is a comparison between new samples and reference samples. This means that new samples can be subjected to the classifier without reenrollment of training data. Only samples which belong to the enrolled class should be positively classified. All other samples are meant to cause a negative authentication result. We can achieve this by training our models to distinguish positive from negative authentications, rather than learning ECG waveforms for every individual. Therefore, difference vectors between feature vectors of different samples are computed by applying a difference function on them. Class patterns are determined by calculating the mean of all windows of the respective class. Class patterns need to be stored to compute difference feature vectors for authentication. Then, for all available windows, difference vectors were calculated and labeled positive, if window and pattern originate from the same class. Otherwise, they are labeled negative. Based on those difference vectors, model training is performed. The trained model is able to distinguish between difference vectors that originate from patterns and samples of the same class and such that originate from patterns and samples from different classes. In case the positive class needs to be exchanged, only the stored class pattern has to be updated. Retraining of the classifier is not involved.
In the next chapter, we thoroughly evaluate the proposed system for both identification and authentication.

Chapter 6

Evaluation

Potential and benefits of continuous ECG classification for authentication and identification have been shown in section 4.4. To assess authentication performance of the system we built and described in chapter 5, we conducted thorough evaluation process. In this chapter, we present details and results of our evaluation.

6.1 FH Hagenberg Research ECG Database

The first step in evaluating our approach is the acquisition of the FH Hagenberg Research ECG Database (FRED). We recorded ECG samples of five minutes from 27 participants, each in one session. 19 participants were male and 8 were female. As shown in figure 6.2, most of our participants were aged from 21 to 25 and 4 participants were from 59 to 62 years old. All participants were apparently healthy and expressed their written consent in capturing, recording, processing and publishing ECG data. There was no

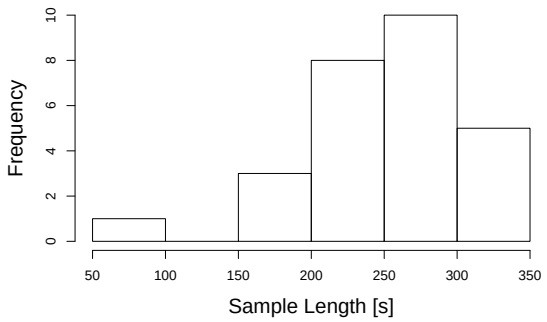

Figure 6.1: Histogram of ECG recording length within FRED.

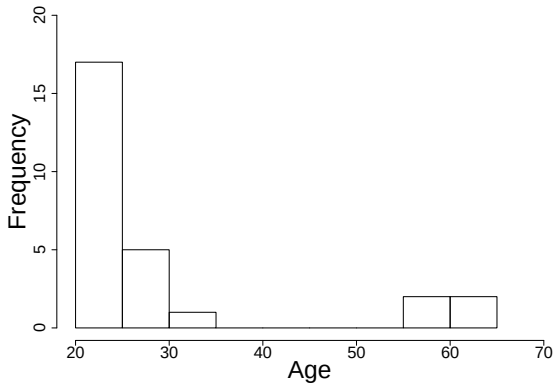

Figure 6.2: Histogram of age of participants, with a median age of 25.

selection of individuals or groups of any kind involved, all participants voluntarily took part in ECG recordings and everyone at our university was invited to participate after a publicly announced invitation.

During data acquisition, the participants were asked to steadily sit on a chair and touch the sensor electrodes with their index fingertips, while keeping a metallic spoon connected to ground in their mouth. Figure 6.3 illustrates our recording setup, while a participant is touching the sensors with his fingertips and keeping a metallic spoon connected to ground in his mouth. Moreover, the participants hands were not supposed to touch each other or connect to the circuitry, to avoid ECG recordings to be biased by short circuits. Movement of the participants and their limbs should be reduced to a minimum, to eliminate interfering electrical charges caused by muscular activity or changes in impedance of the ground electrode.

For some participants, it was difficult or stressful to remain motionless over a period of several minutes. When the ECG recording was massively affected by motion artifacts, biased sections were manually removed before further processing. Figure 6.4 illustrates an ECG signal biased by motion artifacts. Periods of information loss also occurred, when the WIFI signal was temporarily obstructed by other networks or devices. Due to the low transmitting power of our WIFI board, the signal was least obstructed in environments without too many WIFI networks and devices and in close proximity to the access point.

As indicated in figure 6.4, removal of biased sections was performed by cutting at the highest peak of the ECG signal, i.e. R wave, removing the biased section up to the next unbiased peak and merging preceding and subsequent sections. This way, we acquired 24 samples with a minimum length of two minutes. Figure 6.1 depicts the histogram of the length of the acquired raw ECG recordings. 12 of those 24 samples contained about 200 s

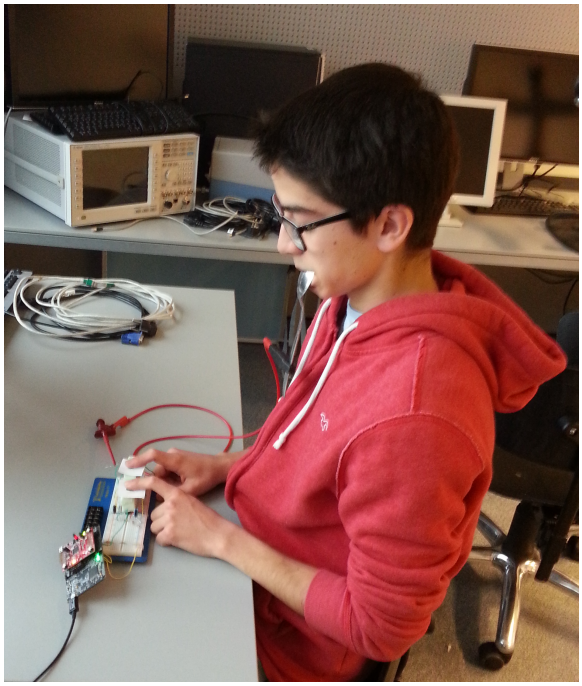

Figure 6.3: ECG recording setup. A participant is sitting on a chair and touching the sensor electrodes, while keeping a spoon connected to ground in his mouth.

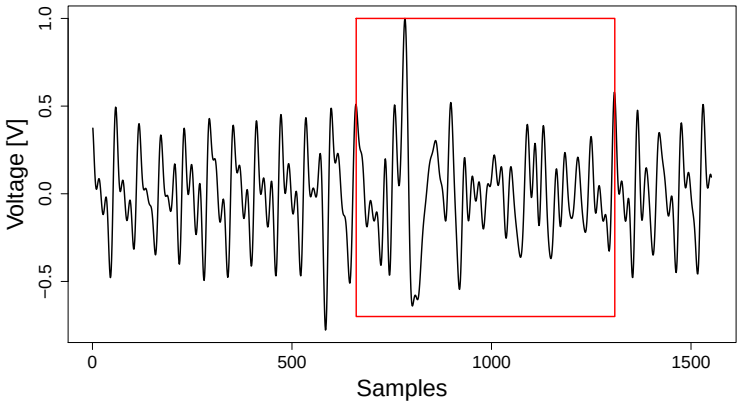

Figure 6.4: ECG recording affected by motion artifact. The highlighted section was manually removed.

of usable data, but in order to keep class balance, sample length is limited by the shortest sample, which has a length of about 110 s. Therefore, every individual in our evaluation is represented by the exact same amount of data.

FRED contains all three datasets: the first dataset contains one raw ECG recording per participant, i.e. 27 recordings of varying length. The second dataset contains 24 samples after digital filtering, manually removing biased sections and cutting recordings to equal length of 110 s. This dataset was selected for evaluation. The third dataset contains 12 samples of equal length of 200 s, preprocessed in an analogous manner.

6.2 Data Partitioning

Our evaluation follows a training/validation/testing scheme. To avoid gallery dependence of our classification models, we initially shuffle our data. Then, a part of our data is separated and forms our held-back test set. This dataset is neither used for training of classifiers, nor to tune hyperparameters for best performance. It is only used for final performance estimation of trained classifiers.

The training data is used for training, validation and parameter tuning of our classifiers. To obtain a realistic performance estimate, we employ cross validation (CV). The caret package [79] for R includes several resampling methods. According to Kim [35], repeated CV outperforms non-repeated CV in obtaining stable performance estimates and therefore can be recommended for general use. Repeated execution on different CV partitions is computationally more expensive than single execution, but for reducing variability of estimators it is worth being carried out [35]. Caret's implementation of repeated CV repeatedly executes K-fold CV, as stated in [38]. Therefore, repeated CV has been used for all classifiers with five repeats and $K = 10$.

6.3 Identification

For identification of participants, each of the 24 ECG recordings contained in FRED was cut into nonoverlapping windows. On each window, autocorrelation was performed. For window length $N = 2.5$ s, each recording results in 44 windows. 24 classes with 44 windows each form a total of 1056 windows that were initially used for classification. 50% of our data has been used for training, while the remaining 50% were used for testing of our classifiers. We commenced evaluation with KNN, SVM and neural network models.

k	Accuracy	Kappa	Accuracy SD	Kappa SD
1	0.363	0.335	0.050	0.052
3	0.378	0.350	0.058	0.060
5	0.395	0.368	0.054	0.057
10	0.402	0.376	0.061	0.063
15	0.415	0.389	0.054	0.057
20	0.395	0.368	0.055	0.058
30	0.383	0.356	0.053	0.056
50	0.361	0.333	0.048	0.050

Table 6.1: Validation parameter grid for KNN ECG subject identification and $N = 2.5\,\mathrm{s}$ for 24 subjects. Best results for $k = 15$.

6.3.1 k-Nearest Neighbors

To receive a first impression of separability of our 24 ECG recordings, we start with simple nearest neighbor classification. To find the number of nearest neighbors, which leads to best adaption of the classifier to our data, parameter grid search is conducted. As shown in table 6.1, $k = 15$ provides best results with an accuracy of 0.41 and Kappa of 0.39 within the tested parameters. Therefore, $k = 15$ is selected for evaluation against the test partition. Figure 6.5a shows the confusion matrix for the classification results of our test partition, which corresponds to an accuracy of 0.38 and Kappa of 0.35. The accuracy is equivalent to the window identification rate (WIR), stated in table 4.1. Although an accuracy of 0.38 seems to be very low, it is well above random performance of $\frac{1}{24}$ within 24 participants. Values along the diagonal from the lower left to upper right corner represent correct classifications, values apart the diagonal represent misclassifications. The confusion matrix indicates a certain discriminability, but the variance is very high. We conclude this from the seemingly random distribution of false classifications over the confusion matrix. The results during cross validation, as shown in table 6.1 are similar to the results achieved from our test partition. We therefore conclude that KNN doesn't suit our data very well and other classifiers than KNN might adapt better.

Before we continue evaluating different classifiers, we try to improve classification performance for KNN. At first we compute the total or subject identification rate. Identification performance can be enhanced, by performing classification on multiple windows of the same subject and determine class membership based on majority voting. Therefore all predictions, i.e. 22 predictions for $N = 2.5\,\mathrm{s}$ and 11 predictions for $N = 5\,\mathrm{s}$ per subject are incorporated resulting in one majority vote each. Figure 6.5b shows the confusion matrix for the same results presented in figure 6.5a, after majority voting was applied. The confusion matrix corresponds to an accuracy of 0.52 and Kappa of 0.5, which is a considerable improvement compared to the WIR of 0.38. Introducing majority voting substantially increases classification performance, but results in longer recording time, as multiple windows

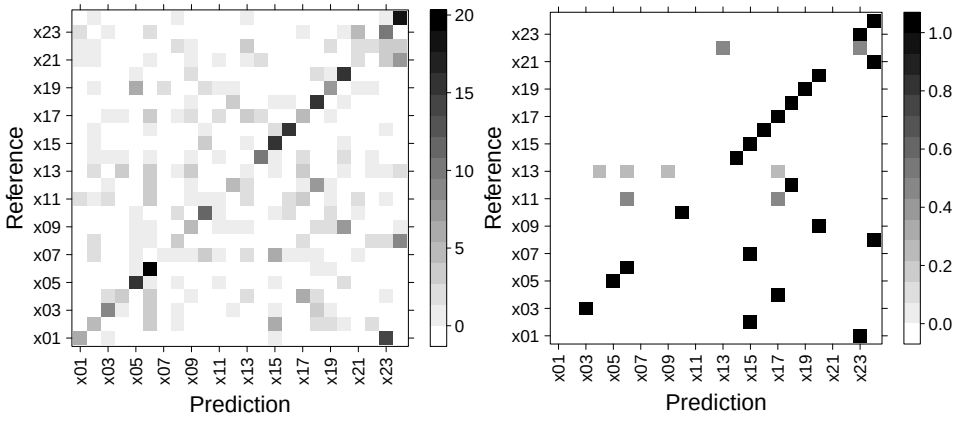

(a) Without majority voting.

(b) With majority voting of 22 predictions.

Figure 6.5: Confusion matrix for ECG subject identification result of test partition for KNN with $k = 15$ and $N = 2.5\,\mathrm{s}$.

k	Accuracy	Kappa	Accuracy SD	Kappa SD
1	0.420	0.394	0.068	0.071
3	0.457	0.433	0.086	0.090
5	0.471	0.448	0.091	0.095
10	0.469	0.445	0.084	0.088
15	0.451	0.427	0.069	0.072
20	0.427	0.402	0.070	0.073
30	0.402	0.376	0.079	0.082
50	0.341	0.313	0.077	0.080

Table 6.2: Validation parameter grid for KNN ECG subject identification and $N = 5\,\mathrm{s}$. Best results for $k = 5$.

need to be recorded, before full authentication performance can be achieved. Therefore we try to reduce variance of our windows by prolonging window length N. Table 6.2 shows the results for our parameter grid for $N = 5\,\mathrm{s}$. The best accuracy and Kappa values were achieved with $k = 5$. On the test partition, this classifier scores an accuracy or WIR of 0.47 and Kappa of 0.45. The corresponding confusion matrix is shown in figure 6.6a.

Figure 6.6b shows the confusion matrix for KNN classification with $k = 5$ and $N = 5\,\mathrm{s}$. It corresponds to an accuracy of 0.72 and Kappa 0.71. Considering this result, we expect more satisfying results from more powerful classifiers.

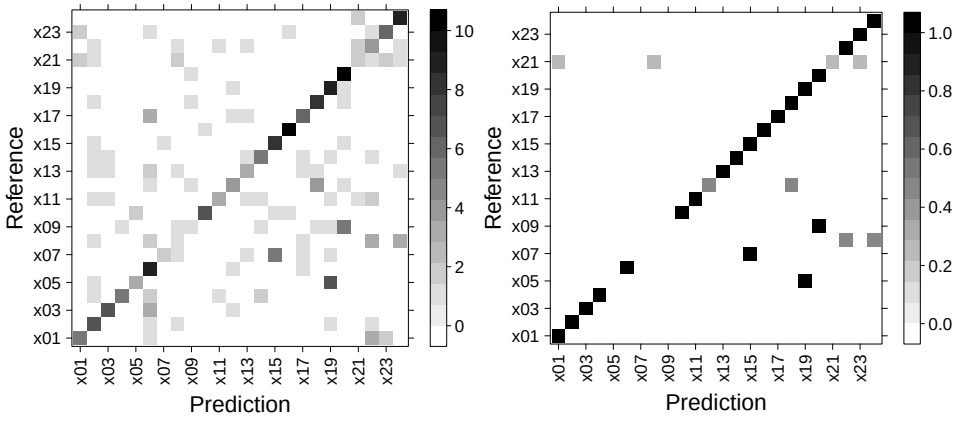

(a) Without majority voting.

(b) With majority voting of 11 predictions.

Figure 6.6: Confusion matrix for ECG subject identification result of test partition for KNN with $k = 5$ and $N = 5\,\text{s}$.

C	Accuracy	Kappa	Accuracy SD	Kappa SD
0.001	0.378	0.354	0.079	0.081
0.01	0.378	0.354	0.079	0.081
0.1	0.494	0.472	0.076	0.079
1	0.490	0.467	0.079	0.081
10	0.450	0.426	0.080	0.084
100	0.448	0.424	0.086	0.089

Table 6.3: Validation parameter grid for linear SVM ECG subject identification and $N = 5\,\text{s}$. Best results for $C = 0.1$.

6.3.2 Support Vector Machine

Analogously to the evaluation process for KNN classification, we conducted parameter grid search for SVMs with linear and radial basis function kernel and for window lengths $N = 2.5\,\text{s}$ and $N = 5\,\text{s}$. Tuning parameter grids for linear and radial SVMs are shown in tables 6.3 and 6.4. As expected, classification with $N = 5\,\text{s}$ outperforms $N = 2.5\,\text{s}$ classification, and majority voting additionally increases accuracy. Confusion matrices for $N = 5\,\text{s}$ are shown in figures 6.7a, 6.7b, 6.8a and 6.8b. Best results were achieved with a radial SVM with $C = 1$ and $\sigma = 0.1$. This classifier achieved a WIR of 0.54 and a SIR of 0.81.

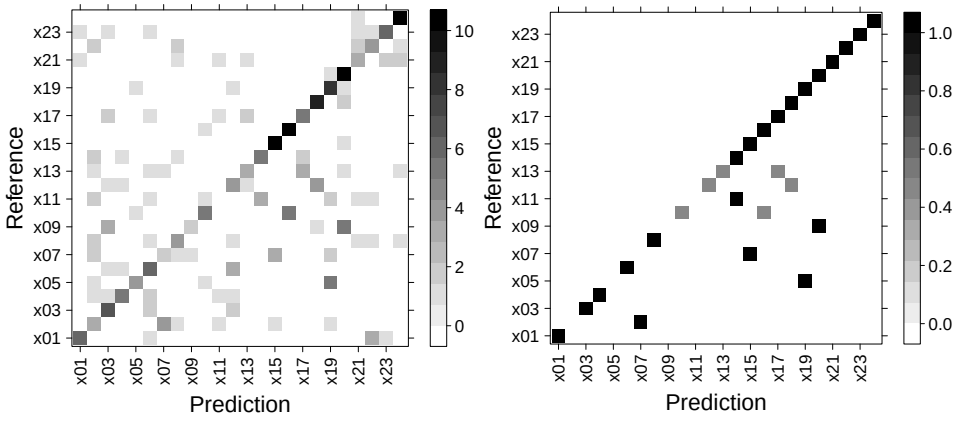

(a) Without majority voting.

(b) With majority voting of 11 predictions.

Figure 6.7: Confusion matrix for linear SVM ECG subject identification on test partition with $C = 0.1$ and $N = 5\,\text{s}$.

σ	C	Accuracy	Kappa	Accuracy SD	Kappa SD
0.01	0.1	0.391	0.368	0.090	0.093
0.01	1.0	0.397	0.373	0.089	0.093
0.01	10.0	0.515	0.494	0.097	0.101
0.01	100.0	0.503	0.480	0.094	0.098
0.10	0.1	0.438	0.415	0.098	0.101
0.10	1.0	0.533	0.513	0.095	0.099
0.10	10.0	0.530	0.509	0.096	0.101
0.10	100.0	0.496	0.473	0.093	0.097
1.00	0.1	0.337	0.311	0.104	0.107
1.00	1.0	0.443	0.419	0.085	0.089
1.00	10.0	0.433	0.408	0.089	0.093
1.00	100.0	0.433	0.408	0.089	0.093

Table 6.4: Validation parameter grid for radial SVM ECG subject identification and $N = 5\,\text{s}$. Best results for $C = 1$, $\sigma = 0.1$.

6.3.3 Neural Network

Additionally, we employed neural networks for classification. Parameter grid search resulted in a parametrization of $size = 30$ and $decay = 1.1$, as shown in table 6.5. Figure 6.9a shows the confusion matrix for neural network classification, figure 6.9b shows the results after majority voting. The achieved WIR of 0.51 and SIR of 0.74 are comparable to our previous results for KNN and SVM. Test accuracy is even slightly better than validation accuracy of 0.5. We therefore believe that the selected parameters lead to an appropriate fit of the model.

Size	Decay	Accuracy	Kappa	Accuracy SD	Kappa SD
1	0.0001	0.085	0.049	0.039	0.039
1	0.001	0.093	0.057	0.044	0.045
1	0.01	0.111	0.076	0.048	0.050
1	0.1	0.082	0.046	0.019	0.020
1	1	0.076	0.039	0.015	0.016
2	0.0001	0.155	0.120	0.062	0.064
2	0.001	0.193	0.159	0.087	0.090
2	0.01	0.198	0.164	0.075	0.077
2	0.1	0.150	0.115	0.052	0.053
2	1	0.144	0.110	0.038	0.039
3	0.0001	0.230	0.197	0.068	0.070
3	0.001	0.234	0.201	0.076	0.079
3	0.01	0.265	0.233	0.076	0.078
3	0.1	0.254	0.222	0.061	0.063
3	1	0.229	0.197	0.060	0.062
5	0.0001	0.331	0.302	0.066	0.069
5	0.001	0.328	0.299	0.086	0.089
5	0.01	0.328	0.298	0.073	0.076
5	0.1	0.360	0.332	0.075	0.077
5	1	0.353	0.325	0.061	0.063
10	0.0001	0.338	0.309	0.085	0.089
10	0.001	0.361	0.333	0.102	0.106
10	0.01	0.373	0.345	0.064	0.067
10	0.1	0.406	0.380	0.086	0.090
10	1	0.454	0.430	0.091	0.095
15	0.0001	0.380	0.352	0.099	0.102
15	0.001	0.379	0.351	0.100	0.104
15	0.01	0.376	0.349	0.094	0.098
15	0.1	0.452	0.427	0.091	0.095
15	1	0.473	0.450	0.095	0.098
20	0.0001	0.385	0.358	0.099	0.102
20	0.001	0.413	0.387	0.105	0.109
20	0.01	0.393	0.366	0.089	0.093
20	0.1	0.440	0.415	0.092	0.096
20	1	0.477	0.453	0.092	0.095
30	0.0001	0.408	0.381	0.098	0.102
30	0.001	0.410	0.384	0.099	0.103
30	0.01	0.423	0.398	0.095	0.099
30	0.1	0.456	0.432	0.088	0.092
30	1	0.485	0.462	0.101	0.105
30	1	0.497	0.475	0.089	0.093
30	1.1	0.504	0.483	0.083	0.086
30	1.2	0.497	0.475	0.083	0.086
30	1.3	0.491	0.468	0.079	0.083
30	1.4	0.490	0.468	0.087	0.091
30	1.5	0.496	0.473	0.089	0.093
30	1.6	0.492	0.470	0.088	0.092
30	1.7	0.493	0.471	0.084	0.088
30	1.8	0.484	0.461	0.086	0.090
30	1.9	0.492	0.469	0.083	0.086
30	2	0.486	0.463	0.086	0.090

Table 6.5: Validation parameter grid for neural network ECG subject identification and $N = 5\,\mathrm{s}$. Best results for $decay = 1.1$, $size = 30$.

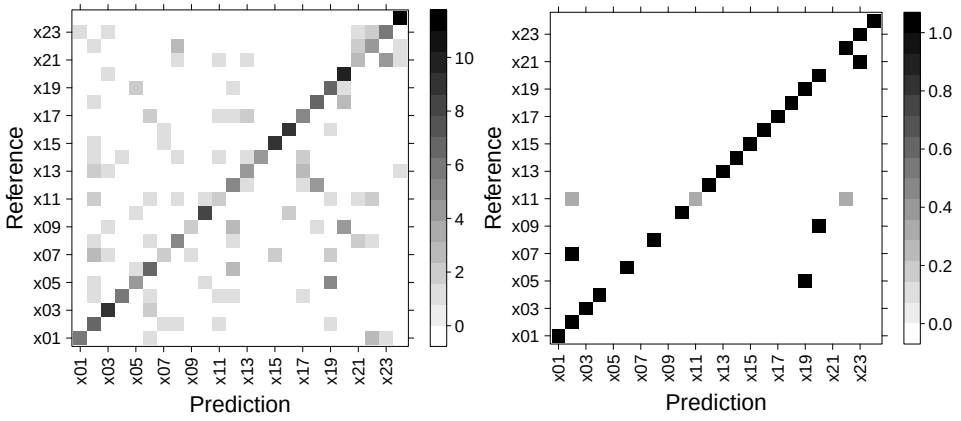

(a) Without majority voting.

(b) With majority voting of 11 predictions.

Figure 6.8: Confusion matrix for radial SVM ECG subject identification on test partition with $C = 1$, $\sigma = 0.1$ and $N = 5\,\mathrm{s}$.

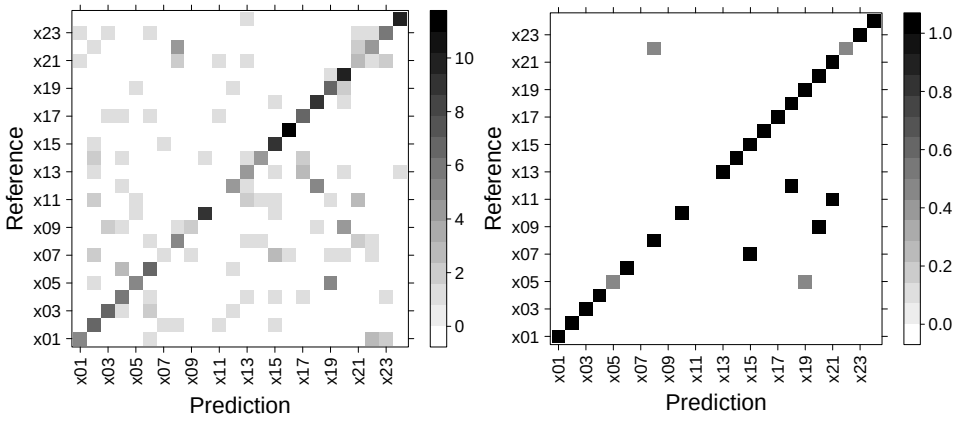

(a) Without majority voting.

(b) With majority voting of 11 predictions.

Figure 6.9: Confusion matrix of neural network ECG subject identification on test partition, $N = 5\,\mathrm{s}, size = 30$ and $decay = 1.1$.

Classifier	WIR	Kappa (WIR)	SIR	Kappa (SIR)
KNN ($k = 5$)	0.47	0.45	0.72	0.71
linear SVM ($C = 0.1$)	0.50	0.47	0.73	0.72
radial SVM ($C = 1, \sigma = 0.1$)	0.54	0.52	0.81	0.80
Neural Network ($size = 30, decay = 1.1$)	0.51	0.49	0.74	0.73

Table 6.6: Classification results for ECG subject identification on test partition, $N = 5\,\mathrm{s}$.

6.3.4 Results

We employed KNN, linear and radial SVMs as well as neural networks for classification. As shown in table 6.6, we received comparable results from different classifiers. We therefore conclude our models to be appropriate and plausible estimations of class separation for our dataset. Best results were achieved with radial SVM classification, $N = 5\,\mathrm{s}$ and majority voting.

6.3.5 Discussion

Our results indicate that ECG has a certain discriminability. Although some classes were correctly predicted with confidence, others were repeatedly misclassified by different classifiers. This could be explained by Doddington's Zoo [31]. Our database seems to contain goats, i.e. classes with high FRR and lambs, i.e. classes with high FAR.

We found that majority voting significantly increases classification performance. Note that our implementation of majority voting is similar to an electoral college. Class membership of all windows of the same subject is assigned to the class with the most votes. If two or more classes have the same amount of votes, classes are assigned in an according split. As many windows contribute to one authentication, the samples per row add up to 1.

Another possibility of increasing classification performance is prolonging window length. It seems to reduce variance upon same class windows, as short term changes in ECG recordings have less impact on longer windows. However, the maximum acceptable window length depends on the use case and is a tradeoff between the desired identification performance and system response time.

Prolonging of window length seems to be better suited for small accuracies and short response times than majority voting, as by doubling window length from $N = 2.5\,\mathrm{s}$ to $N = 5\,\mathrm{s}$, WIR of KNN was raised from 0.38 to 0.52. Using majority voting on two subsequent windows is almost pointless, as the chance for both windows being correctly classified for this case is 0.14, the chance for both windows being incorrectly classified is 0.3844 and for the remaining chance for one correct and one incorrect classification of 0.4712, majority voting could not positively influence the resulting SIR, as this scenario results in a draw. Majority voting therefore is more suitable

for higher classification rates and more windows being incorporated.

However, our system is able to increase performance of existing identification systems by adding confidence in identification results and possibly reducing response time by short circuit evaluation.

6.4 Authentication

Classification for authentication follows a slightly different routine than for identification. From our 24 classes, autocorrelation functions of 528 nonoverlapping windows of size $N = 5\,\mathrm{s}$ are computed. Then, for every class, an ECG pattern is calculated, by taking the mean over all autocorrelated windows from the same class. Afterwards, for all windows, difference vectors between each pattern and window are calculated and stored. Vectors that originate from the same class as the pattern are labeled positive. Vectors derived from classes other than the class of the pattern are labeled negative. This way, 12672 difference vectors are available for classification. For training and validation, 10% of those difference vectors are fed into the classifiers, while 90% are used for our held-back test set. Instead of learning individual ECG waveforms, the models should adapt to pattern-sample difference of positive and negative authentications. In positive cases, pattern and sample are similar, therefore the difference is small. In negative cases, pattern and sample are dissimilar, therefore the difference between them is large.

This procedure implies that our dataset contains 23 negative samples for every positive sample. This imbalance caused our predictors to nearly always predict negative results, while being correct in 23 out of 24 cases. To address this issue, we excluded portions of negative samples from training data. Figure 6.10 depicts the distribution of predicted results for positive and negative samples for different ratios of positive to negative samples (P/N ratio) in the training set. Figure 6.10i shows the original distribution of 1 positive to 23 negative samples. Predicted values for positive and negative samples are hardly seperable. Figures 6.10a to 6.10h depict distributions of predicted results for P/N ratios from 1:1 to 1:8.

A P/N ratio of 1:1 leads to a close distribution of positive samples, as variance within positive samples is smaller than within negative samples. There might be negative samples in the test partition, that differ to negative samples in the training set. Therefore, the predictor adjusts better to the positive than to the negative class. The more negative samples are added to the training set, the more diverse the representation of the negative class gets. Furthermore, with a higher number of negative samples within the training set, positive samples contribute less to overall accuracy. As shown in table 6.7, with more negative samples within the training set, accuracy increases while Kappa decreases. This indicates underrepresentation of positive samples. Although we achieved the highest validation results for a P/N

P/N ratio	Accuracy	Kappa
1:1	0.903	0.806
1:2	0.831	0.613
1:3	0.843	0.560
1:4	0.851	0.516
1:5	0.877	0.535
1:6	0.889	0.467
1:7	0.892	0.331
1:8	0.898	0.200
1:23	0.958	0.000

Table 6.7: Validation results for ECG authentication with different P/N ratios.

k	Accuracy	Kappa	Accuracy SD	Kappa SD
1	0.892	0.490	0.045	0.194
3	0.916	0.564	0.033	0.168
5	0.909	0.522	0.040	0.201
10	0.906	0.517	0.043	0.193
100	0.889	0.000	0.009	0.000

Table 6.8: Validation parameter grid for KNN ECG authentication with regular differences and $N = 5\,\text{s}$. Best results for $k = 3$.

ratio of 1:1, we select 1:2 for further evaluation, as a P/N ratio of 1:1 exposed a gap between validation and test performance. Based on the class separation visible in figure 6.10, we are confident that a P/N ratio of 1:2 fits our data sufficiently.

6.4.1 k-Nearest Neighbors

We start the evaluation process with KNN. Table 6.8 shows our parameter grid for KNN classification. Because of class imbalance in training set, accuracy and Kappa diverge, especially for $k = 100$. Best Kappa and accuracy are achieved with $k = 3$ which is therefore selected for evaluation. As for identification, authentication results are subject to majority voting. For authentication, all difference vectors that originate from one ECG pattern and one class contribute to the vote. Consequently, majority voting significantly increases authentication performance. Figures 6.11a and 6.11b show the ROC for KNN classification prior to and after majority voting. However, to further improve classification performance, we use another difference function. Instead of just calculating the difference, we square the resulting difference vector. Calculating the squared difference introduces weighting, as it emphasizes bigger differences that result from dissimilar ECG windows, in contrast to small differences from similar ECG windows. Table 6.9 shows the parameter grid for KNN with squared differences.
Given the simplicity of this classifier, authentication performance is surprisingly high. With majority voting, ECG authentication with KNN classifi-

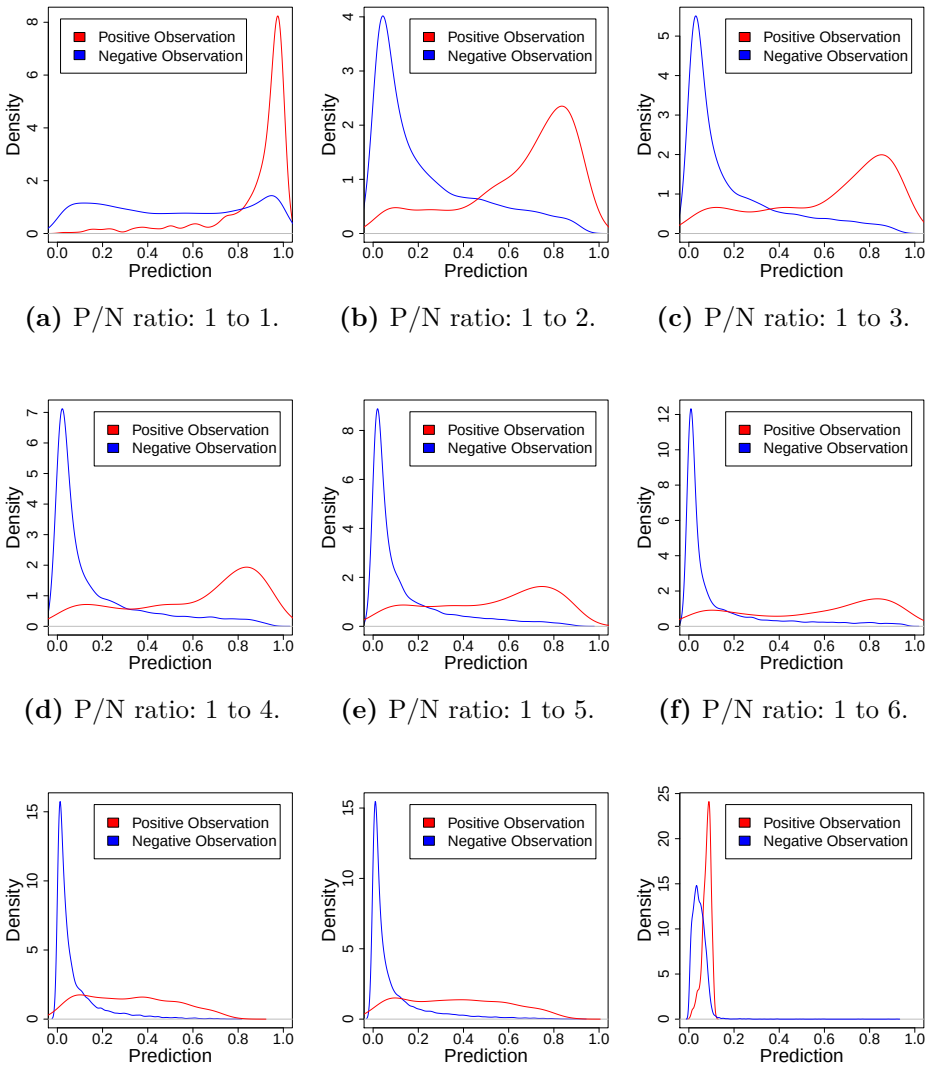

(a) P/N ratio: 1 to 1. **(b)** P/N ratio: 1 to 2. **(c)** P/N ratio: 1 to 3.

(d) P/N ratio: 1 to 4. **(e)** P/N ratio: 1 to 5. **(f)** P/N ratio: 1 to 6.

(g) P/N ratio: 1 to 7. **(h)** P/N ratio: 1 to 8. **(i)** P/N ratio: 1 to 23.

Figure 6.10: Positive and negative class density over predicted value.

k	Accuracy	Kappa	AccuracySD	KappaSD
1	0.776	0.501	0.090	0.200
3	0.790	0.536	0.102	0.219
5	0.818	0.589	0.089	0.194
10	0.824	0.609	0.100	0.217
100	0.667	0.000	0.021	0.000

Table 6.9: Validation parameter grid for KNN ECG authentication with squared differences and $N = 5\,\mathrm{s}$. Best results for $k = 10$.

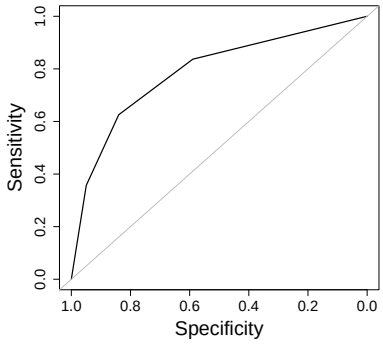

(a) Without majority voting, regular difference. $EER = 0.276, AUC = 0.7873$.

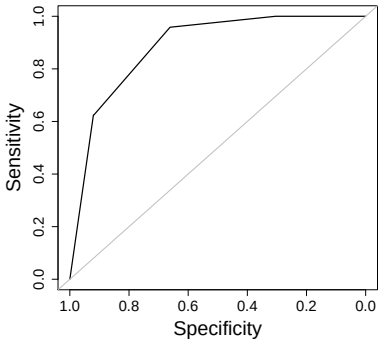

(b) With majority voting, regular difference. $EER = 0.209, AUC = 0.8835$.

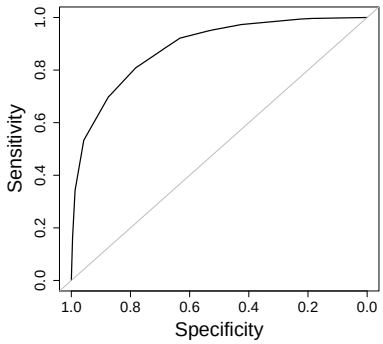

(c) Without majority voting, squared difference. $EER = 0.205, AUC = 0.8840$.

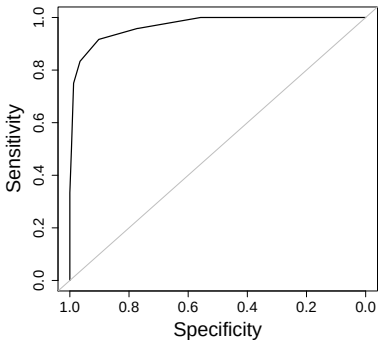

(d) With majority voting, squared difference. $EER = 0.092, AUC = 0.9690$.

Figure 6.11: ROC for KNN ECG authentication on test partition with $k = 3$ for regular difference and $k = 10$ for squared difference. Line segments are caused by small values of k.

cation achieves an EER of 0.092, with an AUC of 0.9690. Figures 6.11c and 6.11d show the ROC for KNN classification with squared differences. Note that line segments are caused by small values of k. For $k = 3$, predicted values lie within $pred \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$.

C	Accuracy	Kappa	Accuracy SD	Kappa SD
0.0001	0.655	0.059	0.060	0.138
0.001	0.661	0.114	0.085	0.214
0.01	0.658	0.127	0.077	0.180
0.1	0.661	0.103	0.090	0.214
1	0.644	0.050	0.091	0.205
10	0.638	0.035	0.082	0.187
100	0.643	0.044	0.083	0.184
1000	0.648	0.063	0.089	0.200
10000	0.657	0.058	0.090	0.224

Table 6.10: Validation parameter grid for linear SVM ECG authentication with regular differences and $N = 5\,\mathrm{s}$. Best results for $C = 0.01$.

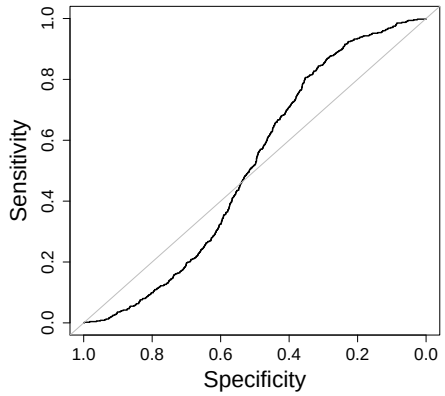

Figure 6.12: ROC for linear SVM ECG authentication on test partition, $C = 0.01$. Data might not be linearly separable. $EER = 0.491, AUC = 0.5154$

6.4.2 Support Vector Machine

Afterward KNN, we trained a linear SVM with our data. Although we tested a variety of parameters from 10^{-4} to 10^4, we were not able to find a suitable model. As shown in table 6.10, for all tested values of C, accuracy and especially Kappa were very low. Figure 6.12 depicts the according ROC, which represents nearly random classification. We conclude that our data might not be linearly separable. When linear SVM classification was applied on squared difference vectors, classes were indeed separable, as shown in table 6.11. This might be due to the fact that simple difference vectors contain both positive and negative values. Squaring removes negative differences and makes it easier for linear models to separate classes. Figure 6.13 depicts ROC for linear SVM classification on squared differences. We continue evaluation with nonlinear classifiers. We subjected our training partition to

C	Accuracy	Kappa	AccuracySD	KappaSD
0.0001	0.757	0.499	0.108	0.208
0.001	0.767	0.522	0.108	0.208
0.01	0.762	0.513	0.110	0.210
0.1	0.835	0.626	0.089	0.196
1	0.848	0.660	0.087	0.186
10	0.840	0.636	0.082	0.187

Table 6.11: Validation parameter grid for linear SVM ECG authentication with squared differences and $N = 5\,\text{s}$. Best results for $C = 1$.

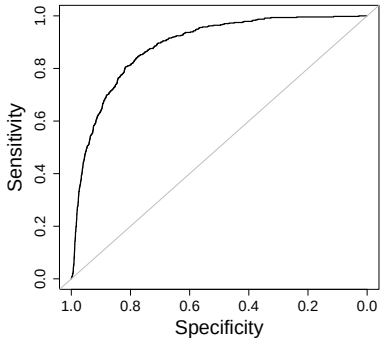

(a) Without majority voting. $EER = 0.191, AUC = 0.8849$.

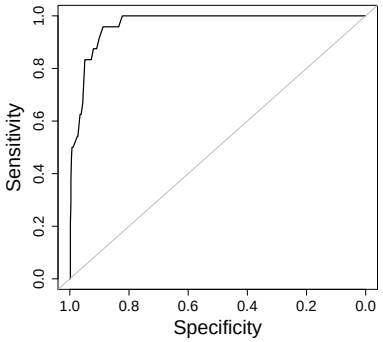

(b) With majority voting. $EER = 0.097, AUC = 0.9669$.

Figure 6.13: ROC for linear SVM ECG authentication on test partition, $C = 1$ for squared difference.

radial SVM classification and achieved best results with $\sigma = 0.1, C = 10$ for regular differences and $\sigma = 0.01, C = 10$ for squared differences. The according parameter grid is depicted in tables 6.12 and 6.13. Classification results for regular distance are shown in figures 6.14a and 6.14b, while squared distance results are depicted in figures 6.14c and 6.14d. Interestingly, considering EER and AUC, results of radial SVM classification are comparable to KNN classification results.

6.4.3 Neural Network

Furthermore, we used neural works for classification. Although the neural network with $size = 15, decay = 0.01$ during validation scored an accuracy value of 0.94 and Kappa of 0.87 as shown in table 6.14, it didn't maintain performance on the test partition. EER of 0.38 and AUC of 0.67 are below average results within our evaluation. However, when applied on squared difference features, a neural network with $size = 1, decay = 1$ achieved the

σ	C	Accuracy	Kappa	Accuracy SD	Kappa SD
0.001	0.001	0.667	0.000	0.022	0.000
0.001	0.01	0.667	0.000	0.022	0.000
0.001	0.1	0.679	0.162	0.080	0.224
0.001	1	0.691	0.174	0.083	0.253
0.001	10	0.679	0.151	0.080	0.229
0.001	100	0.689	0.150	0.083	0.233
0.01	0.001	0.667	0.000	0.022	0.000
0.01	0.01	0.777	0.489	0.103	0.249
0.01	0.1	0.783	0.513	0.101	0.234
0.01	1	0.796	0.543	0.101	0.229
0.01	10	0.836	0.614	0.094	0.235
0.01	100	0.853	0.656	0.086	0.216
0.1	0.001	0.827	0.614	0.095	0.224
0.1	0.01	0.821	0.605	0.100	0.225
0.1	0.1	0.827	0.616	0.101	0.228
0.1	1	0.858	0.665	0.097	0.233
0.1	10	0.872	0.704	0.085	0.204
0.1	100	0.860	0.676	0.073	0.179
1	0.001	0.836	0.580	0.064	0.183
1	0.01	0.841	0.612	0.072	0.186
1	0.1	0.844	0.617	0.073	0.188
1	1	0.836	0.592	0.081	0.217
1	10	0.833	0.589	0.075	0.198
1	100	0.830	0.582	0.076	0.201

Table 6.12: Validation parameter grid for radial SVM ECG authentication with regular differences and $N = 5\,\text{s}$. Best results for $\sigma = 0.1, C = 10$.

σ	C	Accuracy	Kappa	Accuracy SD	Kappa SD
0.001	0.001	0.667	0.000	0.024	0.000
0.001	0.01	0.667	0.000	0.024	0.000
0.001	0.1	0.759	0.500	0.096	0.197
0.001	1	0.756	0.495	0.090	0.182
0.001	10	0.759	0.499	0.088	0.177
0.001	100	0.833	0.618	0.084	0.193
0.01	0.001	0.667	0.000	0.024	0.000
0.01	0.01	0.740	0.461	0.095	0.192
0.01	0.1	0.754	0.491	0.093	0.184
0.01	1	0.748	0.479	0.087	0.172
0.01	10	0.848	0.650	0.085	0.199
0.01	100	0.830	0.612	0.091	0.208
0.1	0.001	0.784	0.530	0.098	0.213
0.1	0.01	0.775	0.514	0.097	0.212
0.1	0.1	0.777	0.515	0.095	0.207
0.1	1	0.824	0.585	0.089	0.213
0.1	10	0.814	0.557	0.095	0.236
0.1	100	0.796	0.508	0.076	0.193
1	0.001	0.819	0.548	0.080	0.217
1	0.01	0.816	0.542	0.081	0.220
1	0.1	0.817	0.545	0.081	0.220
1	1	0.827	0.560	0.072	0.207
1	10	0.817	0.533	0.082	0.234
1	100	0.789	0.451	0.076	0.228

Table 6.13: Validation parameter grid for radial SVM ECG authentication with squared differences and $N = 5\,\text{s}$. Best results for $\sigma = 0.01, C = 10$.

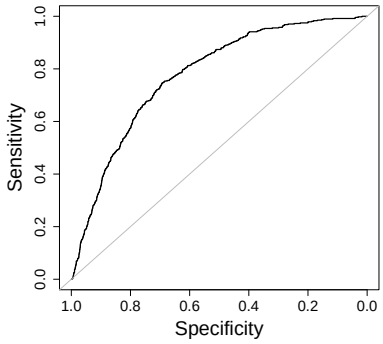

(a) Without majority voting, regular difference. $EER = 0.286, AUC = 0.7752$.

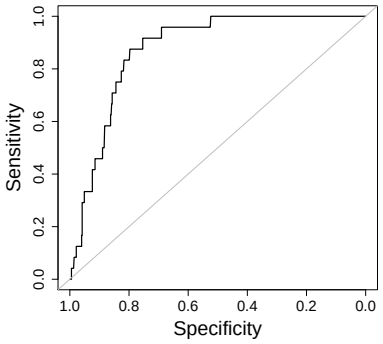

(b) With majority voting, regular difference. $EER = 0.175, AUC = 0.8748$.

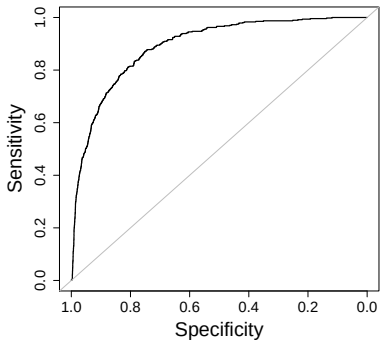

(c) Without majority voting, squared difference. $EER = 0.191, AUC = 0.8904$.

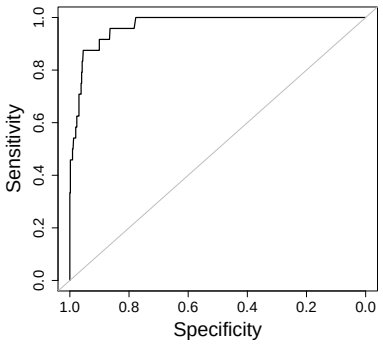

(d) With majority voting, squared difference. $EER = 0.091, AUC = 0.9687$.

Figure 6.14: ROC for radial SVM ECG authentication on test partition, $\sigma = 0.1$ for regular difference, $\sigma = 0.01$ for squared difference, $C = 10$.

highest performance during validation, as shown in table 6.15. On our test partition, achieved an performance of $EER = 0.177, AUC = 0.9024$ and $EER = 0.072, AUC = 0.9811$ with majority voting. Figure 6.15 shows our classification results.

6.4.4 Results

We started evaluation with finding a suitable ratio of negative to positive samples inside the training set. This is important, as too many negative

size	decay	Accuracy	Kappa	AccuracySD	KappaSD
1	0.001	0.77	0.55	0.1	0.185
1	0.01	0.777	0.559	0.098	0.185
1	0.1	0.792	0.58	0.094	0.184
1	1	0.776	0.521	0.107	0.238
3	0.001	0.873	0.718	0.079	0.179
3	0.01	0.892	0.76	0.078	0.172
3	0.1	0.891	0.759	0.074	0.163
3	1	0.883	0.732	0.085	0.198
5	0.001	0.888	0.751	0.063	0.137
5	0.01	0.903	0.787	0.088	0.19
5	0.1	0.921	0.825	0.071	0.152
5	1	0.877	0.721	0.074	0.166
10	0.001	0.91	0.801	0.066	0.147
10	0.01	0.921	0.827	0.065	0.14
10	0.1	0.934	0.854	0.053	0.118
10	1	0.894	0.757	0.069	0.159
15	0.001	0.896	0.779	0.089	0.18
15	0.01	0.94	0.867	0.058	0.131
15	0.1	0.938	0.863	0.054	0.119
15	1	0.893	0.753	0.074	0.173
20	0.001	0.911	0.806	0.068	0.146
20	0.01	0.933	0.85	0.053	0.119
20	0.1	0.938	0.864	0.051	0.11
20	1	0.894	0.757	0.065	0.15

Table 6.14: Validation parameter grid for neural network ECG authentication with regular differences and $N = 5\,\text{s}$. Best results for $size = 15$ and $decay = 0.01$.

samples lead classifiers to mainly adapt to negative samples, while positive samples hardly contribute to the model (figure 6.10h). If the number of negative samples in the training set is too low, classifiers mainly adapt to positive samples, while negative samples are underrepresented (figure 6.10a). Positive samples seem to be more densely populated in feature space, while negative samples are more sparsely distributed over feature space, as they are more heterogeneous. Loss minimization during model training therefore leads the model to adapt to the more homogenous class. In both cases, dominant-class predictions are tightly arranged, while nondominant-class predictions are randomly distributed. For extreme class imbalance (figure 6.10i), predicted results for positive and negative samples are hardly separable. We found a positive/negative sample ratio of 1:2 suitable for our data and predictors.

We used KNN, SVM and neural networks for classification. When applying our data to a linear SVM, we figured out that our data might not be linearly separable. KNN and radial SVM classification provided us with comparable results. Neural network classification scored high validation results, but lowest test results in our validation. We conclude that the model might overfitted our data.

For every classifier, we conducted majority voting. Furthermore, we tested

size	decay	Accuracy	Kappa	AccuracySD	KappaSD
1	0.0001	0.748	0.485	0.104	0.197
1	0.001	0.755	0.496	0.114	0.221
1	0.01	0.730	0.431	0.114	0.238
1	0.1	0.795	0.542	0.112	0.254
1	1	0.812	0.585	0.088	0.193
2	0.0001	0.783	0.534	0.111	0.232
2	0.001	0.771	0.505	0.102	0.220
2	0.01	0.776	0.499	0.096	0.215
2	0.1	0.789	0.527	0.103	0.236
2	1	0.799	0.561	0.090	0.191
3	0.0001	0.804	0.556	0.104	0.243
3	0.001	0.787	0.525	0.104	0.229
3	0.01	0.788	0.528	0.113	0.255
3	0.1	0.789	0.522	0.096	0.214
3	1	0.803	0.571	0.093	0.201
5	0.0001	0.790	0.528	0.093	0.215
5	0.001	0.786	0.515	0.105	0.239
5	0.01	0.799	0.541	0.094	0.232
5	0.1	0.797	0.532	0.097	0.222
5	1	0.801	0.566	0.094	0.199
10	0.0001	0.770	0.470	0.104	0.235
10	0.001	0.774	0.486	0.100	0.227
10	0.01	0.790	0.520	0.099	0.226
10	0.1	0.798	0.533	0.103	0.240
10	1	0.8	0.567	0.094	0.196
15	0.0001	0.774	0.474	0.100	0.226
15	0.001	0.769	0.466	0.096	0.221
15	0.01	0.799	0.532	0.106	0.252
15	0.1	0.797	0.532	0.096	0.217
15	1	0.803	0.571	0.094	0.199
20	0.0001	0.787	0.500	0.096	0.228
20	0.001	0.788	0.506	0.101	0.230
20	0.01	0.794	0.517	0.091	0.221
20	0.1	0.800	0.540	0.106	0.247
20	1	0.810	0.582	0.091	0.198

Table 6.15: Validation parameter grid for neural network ECG authentication with squared differences and $N = 5\,\text{s}$. Best results for $size = 1$ and $decay = 1$.

classification on squared difference vectors. Squared differences emphasize bigger differences compared to smaller ones, which leads to better separation of positive and negative class samples. Furthermore, squaring introduces commutativity and enhanced separability of positive and negative classes. All tested classifiers achieved better results for squared differences than for regular differences. Table 6.16 shows our evaluation results. Best results of EER=0.177 and AUC=0.9024 were achieved with neural network classification on squared differences. Majority voting excels this result with EER=0.072 and AUC=0.9811.

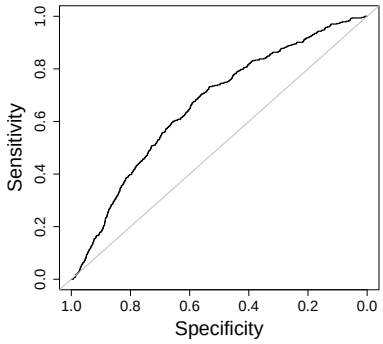

(a) Without majority voting, regular difference. $EER = 0.379, AUC = 0.6717$.

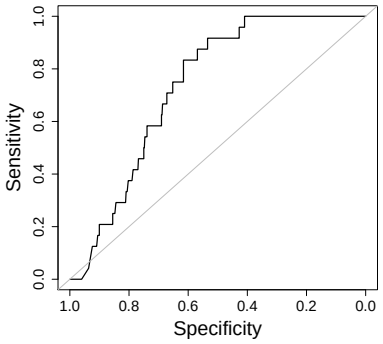

(b) With majority voting, regular difference. $EER = 0.328, AUC = 0.7333$.

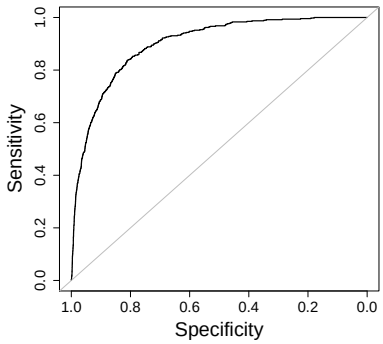

(c) Without majority voting, squared difference. $EER = 0.177, AUC = 0.9024$.

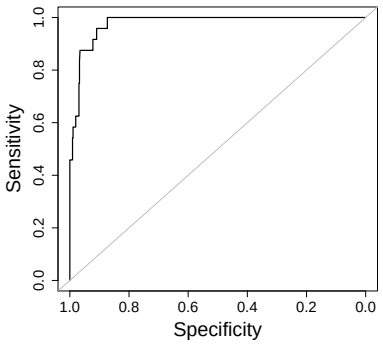

(d) With majority voting, squared difference. $EER = 0.072, AUC = 0.9811$.

Figure 6.15: ROC for neural network ECG authentication on test partition. $size = 15, decay = 0.01$ for regular distance, $size = 20, decay = 1$ for squared difference.

6.4.5 Discussion

The results we achieved for ECG authentication comply with results stated in literature [51]. With 5 s of ECG data, our system provides an EER of 0.177. When used in conjunction with continuous authentication, multiple windows are available for authentication for any point in time after initial authentication. When multiple, consecutive windows are used for authentication based on majority voting, EER as low as 0.072 can be achieved.

Classifier	EER	AUC	EER (maj.V.)	AUC (maj.V.)
KNN ($k = 3$)	0.276	0.7873	0.209	0.8835
KNN squ.diff. ($k = 10$)	0.205	0.8840	0.092	0.9690
linear SVM ($C = 0.01$)	0.491	0.5154	-	-
linear SVM ($C = 1$)	0.191	0.8849	0.097	0.9669
radial SVM ($\sigma = 0.1, C = 10$)	0.286	0.7752	0.175	0.8748
radial SVM squ.diff. ($\sigma = 0.01, C = 10$)	0.191	0.8904	0.091	0.9687
Neural Network ($size = 15, decay = 0.01$)	0.379	0.6717	0.328	0.7333
Neural Network squ.diff. ($size = 1, decay = 1$)	0.177	0.9024	0.072	0.9811

Table 6.16: Authentication results on test partition. Best results for neural network on squared difference vectors.

Therefore, successful authentication over several window periods provides considerable confidence in the result. Based on our findings, we believe that continuous ECG authentication could support increasing mobile device security. Moreover, as ECG data can be recorded unobtrusively, ECG authentication can be combined arbitrarily with different biometrics and systems without necessarily lowering usability.

Although unobtrusively, ECG cannot easily be captured from a distance. The main part of conscious user interaction necessary for continuous ECG authentication is when users attach the ECG recording system, e.g. put on a wearable or smartwatch. It might be difficult and require concerted effort to capture a persons ECG without consent. Therefore, ECG signals are harder to acquire maliciously than other biometrics, e.g. frontal face images. We conclude that continuous ECG authentication provides high levels of robustness against spoofing or counterfeiting attacks. Nonetheless, ECG authentication is not immune to those attacks. Attack scenarios could include data acquisition with capacitive sensors as proposed in [48], which operate through layers of clothing, hidden in a chair-back. However, additionally exploiting ECG for authentication within a multimodal framework adds considerable amounts of effort to possible attacks, without adding additional effort to users. Therefore, ECG authentication is capable of contributing to mobile device security in the future.

Chapter 7

Conclusion

In this thesis, we designed, built and evaluated a system for mobile, continuous ECG authentication. We did a short review on several well known biometrics. We found that for most biometrics, security and usability seem to be inversely correlated. Biometrics which are considered as usable are by trend not very secure, while highly secure biometrics often are not very usable. We believe that the amount and the perceived obtrusiveness of user interaction related with authentication are responsible for user acceptance of authentication systems. Usability and user acceptance are particularly important for continuous authentication systems. Security is often considered as expense rather than as asset. If security measures such as authentication obstruct the intended use of a system, they are unlikely to achieve user acceptance. Therefore, only such biometrics, which can be unobtrusively recorded qualify for continuous use. Typically, behavioral biometrics such as keystroke analysis are employed for continuous authentication, as for their acquisition no dedicated user interaction is required.

We studied the ECG and it's properties which, although the human body is subject to constant change, is surprisingly stable for individuals. After adolescence, the adult ECG waveform remains constant to a certain extent. During the normal aging process, mainly the amplitude of ECG waves decreases. Even the heart rate has not a big impact on the healthy ECG waveform, as the pulse mainly affects the intervals between subsequent heartbeats. Depending on the selected features, ECG authentication systems can therefore be robust against changes in heart rate. However, cardiovascular conditions can cause immediate, drastic and unpredictable changes to the ECG. Such changes would require either online learning or reenrollment of training data for ECG identification. In contrast, our authentication approach would not require retraining to adapt to a changed ECG. It would be sufficient to exchange the ECG pattern for calculation of difference vectors.

In our system design, we discussed different ECG sensor technologies and figured that active, capacitive electrodes are suitable for our system. We built

a prototypic system that connects to our computer via WIFI and transmits the recorded ECG signals. We showed that ECG can be recorded from index fingertips, although we had to accept limitations in usability during data recording. It seems to be necessary to shield sensors and system. Nonetheless, we are convinced about the potential for unobtrusive ECG recording of mobile, continuous ECG authentication. We believe that ECG authentication with the right hardware can be nearly as unobtrusively recorded as other behavioral biometrics. Furthermore, ECG authentication has a wide range of application, as ECG sensors could be included into many mobile devices like smartphones, wearables or cars, as well as infrastructure like building entrances or seatbacks of office chairs. Moreover, continuous ECG authentication systems could provide identity verification for third party applications, devices or services. Such an authentication providing system could be implemented in a wearable and provide authentication for systems within the same personal or body area network.

We recorded the FH Hagenberg Research ECG Database (FRED) for performance evaluation of our system. Several classification models were employed for authentication and identification scenarios. While providing high levels of usability comparable to behavioral biometrics, ECG authentication also provides a considerable degree of security. We employed different classification models and achieved comparable results. We achieved an accuracy of 0.81 and Kappa of 0.80 for identification within 24 individuals using SVM classification. We found that a window length of 5 s is preferable compared to 2.5 s and achieved better results if majority voting was used. During authentication, we were able to achieve an equal error rate of 0.177 with only 5 s of ECG data. For continuous use, we achieved an EER of 0.072 when 110 s of ECG data are incorporated in majority voting.

It seems that continuous ECG authentication has beneficial properties regarding security and usability and therefore is able to add to any system. Using continuous authentication on unobtrusively captured data such as ECG enables us to perform frequent user verification and therefore enhance security, while reducing user interaction related to authentication, which improves usability. This in turn adds to user acceptance, which is crucial to authentication techniques.

However, users, system designers and developers need to be aware, that biometrics in general contain sensitive, personal information. Moreover, ECG contains medical information which in particular deserves protection. We therefore encourage using biometric authentication techniques, but strongly recommend to keep possession and control over biometric templates, e.g. by employing smart cards with biometric match-on-card approaches, as presented by Findling, Hölzl, and Mayrhofer [19]. On a personal level, this can be achieved by authentication frameworks, which process ECG data and provide applications with authentication results, as proposed by Hintze et al. [27]. In conjunction with other biometrics such as gait, voice, face, iris or

fingerprint, ECG recognition is able to form distinctive feature sets, which provide sufficient key space while maintaining usability.

References

Literature

[1] Andrea F. Abate et al. “2D and 3D face recognition: A survey”. In: *Pattern Recognition Letters* 28.14 (2007). Image: Information and Control, pp. 1885–1906. URL: http://www.sciencedirect.com/science/article/pii/S0167865507000189 (cit. on pp. 7, 8).

[2] K. Adamiak, D. Żurek, and K. Ślot. “Liveness detection in remote biometrics based on gaze direction estimation”. In: *Computer Science and Information Systems (FedCSIS), 2015 Federated Conference on*. Sept. 2015, pp. 225–230 (cit. on p. 8).

[3] F. Agrafioti and D. Hatzinakos. “ECG Based Recognition Using Second Order Statistics”. In: *Communication Networks and Services Research Conference, 2008. CNSR 2008. 6th Annual*. May 2008, pp. 82–87 (cit. on pp. 12, 24, 25, 28, 35, 38, 39, 49).

[4] Julio Angulo and Erik Wästlund. “Exploring Touch-Screen Biometrics for User Identification on Smart Phones”. In: *Privacy and Identity Management for Life: 7th IFIP WG 9.2, 9.6/11.7, 11.4, 11.6/PrimeLife International Summer School, Trento, Italy, September 5-9, 2011, Revised Selected Papers*. Ed. by Jan Camenisch et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 130–143. URL: http://dx.doi.org/10.1007/978-3-642-31668-5_10 (cit. on p. 11).

[5] Margit Antal, László Zsolt Szabó, and Izabella László. “Keystroke Dynamics on Android Platform”. In: *Procedia Technology* 19 (2015), pp. 820–826. URL: http://www.sciencedirect.com/science/article/pii/S221201731500119X (cit. on p. 11).

[6] Kiran S. Balagani et al. “On the discriminability of keystroke feature vectors used in fixed text keystroke authentication”. In: *Pattern Recognition Letters* 32.7 (2011), pp. 1070–1080. URL: http://www.sciencedirect.com/science/article/pii/S0167865511000511 (cit. on p. 10).

[7] H. Baltzakis and N. Papamarkos. “A new signature verification technique based on a two-stage neural network classifier”. In: *Engineering Applications of Artificial Intelligence* 14.1 (2001), pp. 95–103. URL: http://www.sciencedirect.com/science/article/pii/S0952197600000646 (cit. on p. 11).

[8] Jean-françois Bonastre et al. “Person Authentication by Voice : A Need For Caution”. In: *in "Proc. Eurospeech'03*. 2003 (cit. on p. 7).

[9] Patrick Bours. “Continuous keystroke dynamics: A different perspective towards biometric evaluation”. In: *Information Security Technical Report* 17.1–2 (2012). Human Factors and Bio-metrics, pp. 36–43. URL: http://www.sciencedirect.com/science/article/pii/S1363412712000027 (cit. on p. 10).

[10] Christina Braz and Jean-Marc Robert. “Security and Usability: The Case of the User Authentication Methods”. In: *Proceedings of the 18th Conference on L'Interaction Homme-Machine*. IHM '06. Montreal, Canada: ACM, 2006, pp. 199–203. URL: http://doi.acm.org/10.1145/1132736.1132768 (cit. on p. 7).

[11] Jeroen Breebaart et al. “Biometric template protection”. In: *Datenschutz und Datensicherheit - DuD* 33.5 (2009), pp. 299–304. URL: http://dx.doi.org/10.1007/s11623-009-0089-0 (cit. on pp. 14, 15).

[12] Lijun Cai, Lei Huang, and Changping Liu. “Person-specific Face Spoofing Detection for Replay Attack Based on Gaze Estimation”. In: *Biometric Recognition: 10th Chinese Conference, CCBR 2015, Tianjin, China, November 13-15, 2015, Proceedings*. Ed. by Jinfeng Yang et al. Cham: Springer International Publishing, 2015, pp. 201–211. URL: http://dx.doi.org/10.1007/978-3-319-25417-3_25 (cit. on p. 8).

[13] N. Carmona et al. “Aging of ECG characteristics over a five year period”. In: *Computing in Cardiology 2013*. Sept. 2013, pp. 1031–1034 (cit. on p. 26).

[14] Yu M. Chi and Gert Cauwenberghs. “Wireless Non-contact EEG/ECG Electrodes for Body Sensor Networks”. In: *Proceedings of the 2010 International Conference on Body Sensor Networks*. BSN '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 297–301. URL: http://dx.doi.org/10.1109/BSN.2010.52 (cit. on pp. 31, 36).

[15] Yu M. Chi, Patrick Ng, and Gert Cauwenberghs. “Wireless Noncontact ECG and EEG Biopotential Sensors”. In: *ACM Trans. Embed. Comput. Syst.* 12.4 (July 2013), 103:1–103:19. URL: http://doi.acm.org/10.1145/2485984.2485991 (cit. on pp. 31, 36).

[16] K. Delac and M. Grgic. “A survey of biometric recognition methods”. In: *Electronics in Marine, 2004. Proceedings Elmar 2004. 46th International Symposium*. June 2004, pp. 184–193 (cit. on pp. 7–10, 13).

[17] I. Deutschmann, P. Nordström, and L. Nilsson. "Continuous Authentication Using Behavioral Biometrics". In: *IT Professional* 15.4 (July 2013), pp. 12–15 (cit. on p. 4).

[18] S.Z. Fatemian and D. Hatzinakos. "A new ECG feature extractor for biometric recognition". In: *16th International Conference on Digital Signal Processing*. July 2009, pp. 1–6 (cit. on pp. 12, 38, 39).

[19] R. D. Findling, M. Hölzl, and R. Mayrhofer. "Mobile Gait Match-on-Card Authentication from Acceleration Data with Offline-Simplified Models". In: *Proceedings of the 14th International Conference on Advances in Mobile Computing and Multimedia (MoMM 2016)*. ACM, 2016 (cit. on p. 76).

[20] Rainhard Dieter Findling. "Pan Shot Face Unlock: Towards Unlocking Personal Mobile Devices using Stereo Vision and Biometric Face Information from multiple Perspectives". Master's thesis. Hagenberg: University of Applied Sciences Upper Austria, Sept. 2013, p. 111 (cit. on p. 8).

[21] M. Frank et al. "Touchalytics: On the Applicability of Touchscreen Input as a Behavioral Biometric for Continuous Authentication". In: *IEEE Transactions on Information Forensics and Security* 8.1 (Jan. 2013), pp. 136–148 (cit. on p. 10).

[22] R. W. Frischholz and A. Werner. "Avoiding replay-attacks in a face recognition system using head-pose estimation". In: *IEEE International Workshop on Analysis and Modeling of Faces and Gestures*. Oct. 2003, pp. 234–235 (cit. on p. 8).

[23] Cristiano Giuffrida et al. "I Sensed It Was You: Authenticating Mobile Users with Sensor-Enhanced Keystroke Dynamics". In: *Detection of Intrusions and Malware, and Vulnerability Assessment: 11th International Conference, DIMVA 2014, Egham, UK, July 10-11, 2014. Proceedings*. Ed. by Sven Dietrich. Cham: Springer International Publishing, 2014, pp. 92–111. URL: http://dx.doi.org/10.1007/978-3-319-08509-8_6 (cit. on p. 11).

[24] M. Guennoun et al. "Continuous authentication by electrocardiogram data". In: *Science and Technology for Humanity (TIC-STH), 2009 IEEE Toronto International Conference*. Sept. 2009, pp. 40–42 (cit. on pp. 4, 12).

[25] Irene Guggenmoos-Holzmann. "The meaning of kappa: Probabilistic concepts of reliability and validity revisited". In: *Journal of Clinical Epidemiology* 49.7 (1996), pp. 775–782. URL: http://www.sciencedirect.com/science/article/pii/089543569600011X (cit. on p. 50).

[26] Marian Harbach et al. “It’s a Hard Lock Life: A Field Study of Smartphone (Un)Locking Behavior and Risk Perception”. In: *Symposium On Usable Privacy and Security (SOUPS 2014)*. Menlo Park, CA: USENIX Association, July 2014, pp. 213–230. URL: https://www.usenix.org/conference/soups2014/proceedings/presentation/harbach (cit. on pp. 5, 14).

[27] Daniel Hintze et al. “Confidence and Risk Estimation Plugins for Multi-Modal Authentication on Mobile Devices Using CORMORANT”. In: *Proceedings of the 13th International Conference on Advances in Mobile Computing and Multimedia*. MoMM 2015. Brussels, Belgium: ACM, 2015, pp. 384–388. URL: http://doi.acm.org/10.1145/2837126.2843845 (cit. on p. 76).

[28] ISO. *Information technology - Security techniques - Biometric information protection*. Standard ISO/IEC 24745:2011. International Organization for Standardization, 2011 (cit. on p. 15).

[29] Steven A. Israel et al. “ECG to identify individuals”. In: *Pattern Recognition* 38.1 (2005), pp. 133–142. URL: http://www.sciencedirect.com/science/article/pii/S0031320304002419 (cit. on pp. 12, 34, 37, 39).

[30] Anil K. Jain, Karthik Nandakumar, and Abhishek Nagar. “Biometric Template Security”. In: *EURASIP J. Adv. Signal Process* 2008 (Jan. 2008), 113:1–113:17. URL: http://dx.doi.org/10.1155/2008/579416 (cit. on p. 15).

[31] Anil K. Jain, Arun A. Ross, and Karthik Nandakumar. *Introduction to Biometrics*. Springer Publishing Company, Incorporated, 2011 (cit. on p. 62).

[32] A. J. Jerri. “The Shannon sampling theorem - Its various extensions and applications: A tutorial review”. In: *Proceedings of the IEEE* 65.11 (Nov. 1977), pp. 1565–1596 (cit. on p. 47).

[33] I. T. Jolliffe. “Principal Component Analysis and Factor Analysis”. In: *Principal Component Analysis*. New York, NY: Springer New York, 2002, pp. 150–166. URL: http://dx.doi.org/10.1007/0-387-22440-8_7 (cit. on p. 22).

[34] Rupali Khane, Anil D. Surdi, and Rajamati Shakar Bhatkar. “Changes in ECG pattern with advancing age”. In: *Journal of Basic and Clinical Physiology and Pharmacology* 22 (Dec. 2011), pp. 97–101 (cit. on p. 26).

[35] Ji-Hyun Kim. “Estimating classification error rate: Repeated cross-validation, repeated hold-out and bootstrap”. In: *Computational Statistics & Data Analysis* 53.11 (2009), pp. 3735–3745. URL: http://www.sciencedirect.com/science/article/pii/S0167947309001601 (cit. on p. 55).

[36] André Knörig, Reto Wettach, and Jonathan Cohen. "Fritzing: A Tool for Advancing Electronic Prototyping for Designers". In: *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*. TEI '09. Cambridge, United Kingdom: ACM, 2009, pp. 351–358. URL: http://doi.acm.org/10.1145/1517664.1517735 (cit. on pp. 42, 43).

[37] Ulrike Korte et al. "A cryptographic biometric authentication system based on genetic fingerprints". In: *SICHERHEIT 2008. Sicherheit, Schutz und Zuverlässigkeit. Beiträge der 4. Jahrestagung des Fachbereichs Sicherheit der Gesellschaft für Informatik e.V. (GI)*. Ed. by Ammar Alkassar and Jörg Siekmann. 2008, pp. 263–276 (cit. on pp. 11, 12).

[38] Max Kuhn. *A Short Introduction to the caret Package*. 2016. URL: https://cran.r-project.org/web/packages/caret/vignettes/caret.pdf (visited on 11/17/2016) (cit. on p. 55).

[39] Max Kuhn and Kjell Johnson. *Applied predictive modeling*. Springer Science+Business Media, 2013 (cit. on pp. 18–20).

[40] R.D. Labati, R. Sassi, and F. Scotti. "ECG biometric recognition: Permanence analysis of QRS signals for 24 hours continuous authentication". In: *IEEE International Workshop on Information Forensics and Security (WIFS)*. Nov. 2013, pp. 31–36 (cit. on p. 12).

[41] L. Lee and W. E. L. Grimson. "Gait analysis for recognition and classification". In: *Fifth IEEE International Conference on Automatic Face and Gesture Recognition*. May 2002, pp. 148–155 (cit. on p. 9).

[42] M. Li and X. Li. "Verification based ECG biometrics with cardiac irregular conditions using heartbeat level and segment level information fusion". In: *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. May 2014, pp. 3769–3773 (cit. on p. 28).

[43] Lewis A. Lipsitz and Ary L. Goldberger. "Loss of complexity and aging. Potential applications of fractals and chaos theory to senescence". In: *JAMA* (1992), pp. 1806–1809 (cit. on p. 26).

[44] André Lourenço, Hugo Silva, and Ana Fred. "Unveiling the Biometric Potential of Finger-based ECG Signals". In: *Intell. Neuroscience* 2011 (Jan. 2011), 5:1–5:8. URL: http://dx.doi.org/10.1155/2011/720971 (cit. on pp. 12, 30, 36, 38, 39).

[45] C. Maple and P. Norrington. "The usability and practicality of biometric authentication in the workplace". In: *First International Conference on Availability, Reliability and Security (ARES'06)*. Apr. 2006 (cit. on pp. 10, 13).

[46] R. Matias et al. “ECG monitoring via Capacitive Body Coupled Communications”. In: *IEEE International Symposium on Medical Measurements and Applications (MeMeA)*. June 2014, pp. 1–6 (cit. on p. 29).

[47] Warren S. McCulloch and Walter Pitts. “A logical calculus of the ideas immanent in nervous activity”. In: *The bulletin of mathematical biophysics* 5.4 (1943), pp. 115–133. URL: http://dx.doi.org/10.1007/BF02478259 (cit. on p. 19).

[48] N.J. McDonald et al. “Noncontact ECG system for unobtrusive long-term monitoring”. In: *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*. Aug. 2012, pp. 1614–1618 (cit. on pp. 31, 74).

[49] Muhammad Muaaz and René Mayrhofer. “An Analysis of Different Approaches to Gait Recognition Using Cell Phone Based Accelerometers”. In: *Proceedings of International Conference on Advances in Mobile Computing & Multimedia*. MoMM ’13. Vienna, Austria: ACM, 2013, 293:293–293:300. URL: http://doi.acm.org/10.1145/2536853.2536895 (cit. on p. 9).

[50] E. Nemati, M.J. Deen, and T. Mondal. “A wireless wearable ECG sensor for long-term applications”. In: *Communications Magazine, IEEE* 50.1 (Jan. 2012), pp. 36–43 (cit. on p. 31).

[51] I. Odinaka et al. “ECG Biometric Recognition: A Comparative Analysis”. In: *IEEE Transactions on Information Forensics and Security* 7.6 (Dec. 2012), pp. 1812–1824 (cit. on pp. 12, 14, 34, 40, 73).

[52] L. O’Gorman. “Comparing passwords, tokens, and biometrics for user authentication”. In: *Proceedings of the IEEE* 91.12 (Dec. 2003), pp. 2021–2040 (cit. on p. 13).

[53] Chulsung Park et al. “An ultra-wearable, wireless, low power ECG monitoring system”. In: *Biomedical Circuits and Systems Conference, 2006. BioCAS 2006. IEEE*. Nov. 2006, pp. 241–244 (cit. on p. 31).

[54] D. A. Reid et al. “Soft biometrics for surveillance: An overview”. In: *Handbook of Statistics* 31 (2013), pp. 327–352 (cit. on p. 12).

[55] A Searle and L Kirkup. “A direct comparison of wet, dry and insulating bioelectric recording electrodes”. In: *Physiological Measurement* 21.2 (2000), p. 271. URL: http://stacks.iop.org/0967-3334/21/i=2/a=307 (cit. on pp. 29, 31, 32).

[56] T.W. Shen, W.J. Tompkins, and Y.H. Hu. “One-lead ECG for identity verification”. In: *Engineering in Medicine and Biology, 2002. 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society EMBS/BMES Conference, 2002. Proceedings of the Second Joint*. Vol. 1. 2002, 62–63 vol.1 (cit. on pp. 12, 37, 39).

[57] H.P. da Silva et al. “Finger ECG signal for user authentication: Usability and performance”. In: *Sixth IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS)*. Sept. 2013, pp. 1–8 (cit. on pp. 12, 29, 36, 37, 39).

[58] Y. N. Singh and P. Gupta. “ECG to Individual Identification”. In: *2nd IEEE International Conference on Biometrics: Theory, Applications and Systems, BTAS*. Sept. 2008, pp. 1–8 (cit. on p. 26).

[59] L. Sirovich and M. Kirby. “Low-dimensional procedure for the characterization of human faces”. In: *J. Opt. Soc. Am. A* 4.3 (Mar. 1987), pp. 519–524. URL: http://josaa.osa.org/abstract.cfm?URI=josaa-4-3-519 (cit. on p. 8).

[60] Joachim Taelman et al. “Influence of Mental Stress on Heart Rate and Heart Rate Variability”. In: *4th European Conference of the International Federation for Medical and Biological Engineering: ECIFMBE 2008 23–27 November 2008 Antwerp, Belgium*. Ed. by Jos Vander Sloten et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 1366–1369. URL: http://dx.doi.org/10.1007/978-3-540-89208-3_324 (cit. on pp. 27, 28).

[61] Matthew Turk and Alex Pentland. “Eigenfaces for Recognition”. In: *J. Cognitive Neuroscience* 3.1 (Jan. 1991), pp. 71–86. URL: http://dx.doi.org/10.1162/jocn.1991.3.1.71 (cit. on p. 8).

[62] J.A. Unar, Woo Chaw Seng, and Almas Abbasi. “A review of biometric technology along with trends and prospects”. In: *Pattern Recognition* 47.8 (2014), pp. 2673–2688. URL: http://www.sciencedirect.com/science/article/pii/S003132031400034X (cit. on pp. 6–10, 12, 13).

[63] A. C. Weaver. “Biometric authentication”. In: *Computer* 39.2 (Feb. 2006), pp. 96–97 (cit. on p. 9).

[64] Greg Whyte and Sanjay Sharma. *Practical ECG for Exercise Science and Sports Medicine*. Human Kinetics, 2010 (cit. on p. 27).

[65] J. Yoo et al. “A Wearable ECG Acquisition System With Compact Planar-Fashionable Circuit Board-Based Shirt”. In: *IEEE Transactions on Information Technology in Biomedicine* 13.6 (Nov. 2009), pp. 897–902 (cit. on pp. 29, 36).

[66] W. Zhao et al. “Face Recognition: A Literature Survey”. In: *ACM Comput. Surv.* 35.4 (Dec. 2003), pp. 399–458. URL: http://doi.acm.org/10.1145/954339.954342 (cit. on p. 8).

Online sources

[67] Arduino. *Arduino Software*. 2016. URL: https://www.arduino.cc/en/Main/Software (visited on 10/25/2016) (cit. on p. 45).

[68] Anthony Atkielski. *SinusRhythmLabels*. 2016. URL: https://commons.wikimedia.org/w/index.php?title=File:SinusRhythmLabels.svg&oldid=194734208 (visited on 08/23/2016) (cit. on p. 25).

[69] Wikipedia Contributors. *Adermatoglyphia*. 2016. URL: https://en.wikipedia.org/w/index.php?title=Adermatoglyphia&oldid=739457905 (visited on 10/01/2016) (cit. on p. 6).

[70] Wikipedia Contributors. *List of data breaches*. 2016. URL: https://en.wikipedia.org/w/index.php?title=List_of_data_breaches&oldid=743031616 (visited on 10/13/2016) (cit. on p. 14).

[71] Wikipedia Contributors. *Phonology*. 2016. URL: https://en.wikipedia.org/w/index.php?title=Phonology&oldid=738814514 (visited on 10/03/2016) (cit. on p. 7).

[72] Mouser Electronics. *Plessey Semiconductors PS25101*. 2016. URL: http://www.mouser.at/ProductDetail/Plessey-Semiconductors/PS25101/?qs=sGAEpiMZZMt1hNLluIiMiw6K4N52lRD46wMWCZTCqr0= (visited on 11/02/2016) (cit. on p. 45).

[73] Energia. *Energia Open-Source Electronics Prototyping Platform*. 2016. URL: http://energia.nu/ (visited on 10/25/2016) (cit. on pp. 44, 45).

[74] Mikael Häggström. *Precordial leads in ECG*. 2012. URL: https://commons.wikimedia.org/wiki/File:Precordial_leads_in_ECG.png (visited on 11/27/2016) (cit. on p. 30).

[75] Texas Instruments Incorporated. *INA 122 Single Supply, MicroPower INSTRUMENTATION AMPLIFIER*. 2016. URL: http://www.ti.com/lit/ds/symlink/ina122.pdf (visited on 10/24/2016) (cit. on p. 44).

[76] Texas Instruments Incorporated. *MSP432P401R LaunchPad*. 2016. URL: http://www.ti.com/tool/msp-exp432p401r (visited on 10/25/2016) (cit. on p. 44).

[77] Texas Instruments Incorporated. *SimpleLink™ Wi-Fi® CC3100 wireless network processor BoosterPack™ plug-in module*. 2016. URL: http://www.ti.com/tool/cc3100boost (visited on 10/25/2016) (cit. on p. 45).

[78] The MathWorks Incorporated. *Using Filter Designer*. 2016. URL: https://de.mathworks.com/help/dsp/ug/opening-fdatool.html?searchHighlight=fdatool (visited on 10/27/2016) (cit. on pp. 47, 48).

[79] Max Kuhn. *caret: Classification and Regression Training*. 2016. URL: https://cran.r-project.org/package=caret (visited on 11/14/2016) (cit. on p. 55).

[80] Rhcastilhos. *Schematic diagram of the human eye*. 2016. URL: https://commons.wikimedia.org/w/index.php?title=File:Schematic_diagram_of_the_human_eye_horizontal_pt.svg&oldid=199811256 (visited on 10/24/2016) (cit. on p. 9).

[81] Plessey Semiconductors Ltd. *Application Note # 291474 ECG sensor in a SmartPhone*. 2016. URL: http://www.plesseysemiconductors.com/doc/?id=291474 (visited on 10/24/2016) (cit. on p. 43).

[82] Plessey Semiconductors Ltd. *Application Note # 291491 Single arm ECG measurement using EPIC*. 2016. URL: http://www.plesseysemiconductors.com/doc/?id=291491 (visited on 10/24/2016) (cit. on p. 43).

[83] Plessey Semiconductors Ltd. *PS25201A / B EPIC Ultra High Impedance Electrophysiological Sensor*. 2016. URL: http://www.plesseysemiconductors.com/doc/?id=291841 (visited on 10/24/2016) (cit. on pp. 42, 44).

[84] European Union. *Integration of biometric features in passports and travel documents*. 2016. URL: http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:l14154 (visited on 10/15/2016) (cit. on p. 6).

[85] Unknown. *EKG-Aufzeichnung*. 2016. URL: https://de.wikipedia.org/w/index.php?title=Datei:EKG-Aufzeichnung.svg&oldid=143956654 (visited on 08/26/2016) (cit. on p. 25).

[86] Frank G. Yanowitz. *Characteristics of the Normal ECG*. 2016. URL: http://ecg.utah.edu/lesson/3 (visited on 08/23/2016) (cit. on p. 24).

Messbox zur Druckkontrolle

— Druckgröße kontrollieren! —

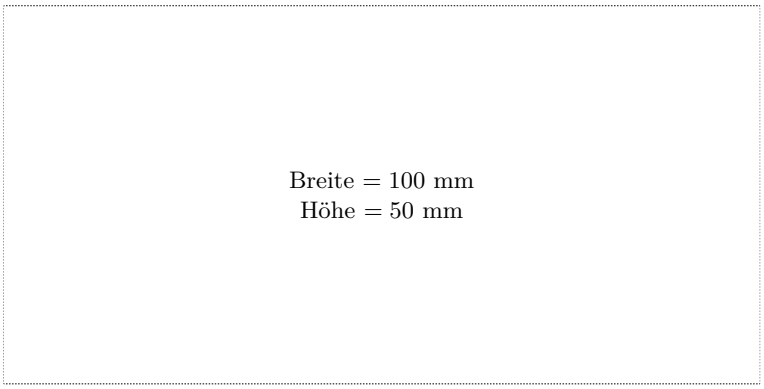

— Diese Seite nach dem Druck entfernen! —