Pulse-Response: Exploring Human Body Impedance for Biometric Recognition

IVAN MARTINOVIC, University of Oxford
KASPER B. RASMUSSEN, University of Oxford
MARC ROESCHLIN, University of Oxford
GENE TSUDIK, University of California, Irvine

Biometric characteristics are often used as a supplementary component in user authentication and identification schemes. Many biometric traits, both physiological and behavioral, offering a wider range of security and stability have been explored. We propose a new physiological trait based on the human body's electrical response to a square pulse signal, called *pulse-response*, and analyze how this biometric characteristic can be used to enhance security in the context of two example applications: (1) an additional authentication mechanism in PIN entry systems, and (2) a means of continuous authentication on a secure terminal. The pulse-response biometric recognition is effective because each human body exhibits a unique response to a signal pulse applied at the palm of one hand, and measured at the palm of the other. This identification mechanism integrates well with other established methods and could offer an additional layer of security, either on a continuous basis or at login time. We build a proof-of-concept prototype and perform experiments to assess the feasibility of pulse-response for biometric authentication. The results are very encouraging, achieving an equal error rate of 2% over a static data set, and 9% over a data set with samples taken over several weeks. We also quantize resistance to attack by estimating individual worst-case probabilities for zero-effort impersonation in different experiments.

CCS Concepts:•**Security and privacy** → **Biometrics;** *Multi-factor authentication; Access control;*

General Terms: Design, Security, Performance

Additional Key Words and Phrases: Access control, Biometric authentication, Biometrics, Biometric recognition, Continuous authentication, User identification, Physiological trait, Secure computer terminal, Secure man-machine interaction, Secure PIN entry

ACM Reference Format:
Ivan Martinovic, Kasper B. Rasmussen, Marc Roeschlin, and Gene Tsudik, 2017. Pulse-Response: Exploring Human Body Impedance for Biometric Recognition. *ACM Trans. Info. Syst. Sec.* 0, 0, Article 0 (0000), 30 pages.
DOI: 0000000.0000000

1. INTRODUCTION

Many modern access control systems augment traditional two-factor authentication (something you know and something you have) with a third factor: "something you are", i.e., biometric authentication. This additional layer of security comes in many flavors: from fingerprint readers on laptops used to facilitate easy login with a single finger swipe, to iris scanners used as auxiliary authentication for accessing secure facilities. In the latter case, the authorized user typically presents a smart card, then types in a PIN, and finally performs an iris (or fingerprint) scan.

In this paper, we propose a new biometric characteristic based on the human body's electrical response to a square pulse signal. We consider two motivating scenarios:

Author's addresses: I. Martinovic and K. B. Rasmussen and M. Roeschlin, Department of Computer Science, University of Oxford, UK; G. Tsudik, Department of Computer Science, University of California, Irvine, California, USA.

DOI: 0000000.0000000

The first is an access control setting where this new biometric mode is used as an additional layer of security when a user enters a PIN, e.g., into a bank ATM. Pulse-response biometric recognition (pulse-response recognition in short) facilitates unification of PIN entry and biometric capture. We use PIN entry as a running example for this scenario throughout the paper. This is because PIN pads (e.g., in ATMs) are often made of metal, which makes capturing the biometric trait straightforward: a user would place one hand on a metal pad adjacent to the key-pad, while using the other hand to enter a PIN. This conductive pad would transmit the pulse and a sensor in the PIN pad would capture the measurement.

The second scenario corresponds to *continuous authentication* at a stationary computer terminal, e.g., verifying that the user, who logged in earlier, is the same person currently at the keyboard. For this scenario, we need a mechanism that periodically samples one or more biometric characteristics. However, for obvious usability reasons, this should ideally be done *unobtrusively*. Pulse-response recognition is particularly well-suited for this setting. Assuming that it can be made from a conductive material, the keyboard would generate the pulse signal and measure the electric response, while the user (remaining oblivious) is typing. The main idea is that the user's pulse-response is captured at login time and identity of the person currently at the keyboard can be verified transparently, at desired frequency.

The continuous authentication problem can be challenging to solve using static biometric modalities. For example, if swipe fingerprint sensors are used, the user of such an authentication system would have to periodically stop and swipe a finger on the scanner, which could be disruptive. Less obtrusive approaches try to solve this problem using automated video monitoring or continuous face recognition [Niinuma and Jain 2010; Sensible Vision Inc. 2013] with a video camera. However, depending on the context, such systems can be perceived as invasive. Unlike many static traits, behavioural characteristics can allow for very non-invasive continuous authentication, most notably keystroke timings and mouse dynamics [Banerjee and Woodard 2012]. By continuously measuring and quantizing the interaction with mouse and keyboard one can verify if an originally logged in user is still present at the computer terminal or if someone else took over an open session. We present a possible solution to continuous authentication that is equally transparent and unobtrusive as a keystroke dynamics but is based on a physiological trait.

To assess efficacy and feasibility of pulse-response recognition, we built a prototype platform for gathering pulse-response data. Its main purpose is to assess whether we can identify users from a population of test subjects. The same platform can test the distinguishing power and stability of this trait over time. We also explored two hypothetical systems that apply pulse-response recognition to the two sample scenarios discussed above: one to unobtrusively capture the biometric characteristic for an additional layer of security when entering a PIN, and the other to implement continuous authentication.

This paper is based on and extends the publication "Authentication Using Pulse-Response Biometrics" [Rasmussen et al. 2014]. In particular, it contains a more comprehensive analysis of pulse-response recognition: Besides average performance of the biometric recognition method in terms of algorithm and system errors, we examine experimentally for each subject in our data set how resistant the biometric trait is to impersonation attacks. To this end, we design experiments with selected combinations of attackers and victims from our test subject population and calculate worst-case probabilities for impersonation. We differentiate between internal and external impostors, i.e., attackers whose biometric template is known to the classification algorithm and attackers who are unknown. In addition, the underlying test subject population of this

work is increased threefold over the publication in [Rasmussen et al. 2014] to achieve higher statistical significance of all reported results.

The rest of the paper is organized as follows: Section 2 provides some background on biometric recognition and states our goals and requirements for the presented biometric modality. Section 3 describes pulse-response recognition in detail. Sections 4 and 5 present the PIN entry and continuous authentication systems, respectively. Section 6 describes the biometric data capture setup and Section 7 presents experimental results. In Section 8 resistance to impersonation is investigated. Related work is overviewed in Section 9 and the paper concludes with Section 10.

2. BACKGROUND

This section provides background on biometric recognition, summarizes the terminology and introduces our design goals.

2.1. Biometric Verification and Identification

Given basic familiarity with the subject, this section can be skipped with no loss of continuity.

The US National Institute of Standards and Technology (NIST) divides biometric measurements into two categories [Information Technology Laboratory – National Institute of Standards and Technology 2013], physiological and behavioral. The former relies on the physiology of a person and includes: fingerprints, hand geometry, facial recognition, speech analysis, and iris/retina scans. Behavioral traits are based on user behavior and include, for instance, keystroke timings, speech pattern analysis or gait recognition and analysis of stylus pressure, acceleration and shape in hand-writing.

Physiological biometric characteristics can help identify an individual among a large pool of candidates. In general, physiological biometrics are considered moderately difficult to circumvent. For example, although hand geometry is very stable over the course of one's adult life, it does not provide enough distinguishing power to be used as the only means for identification [National Science & Technology Council 2006]. Also, facial recognition systems that do not employ liveness detection can be fooled by an appropriately-sized photo of a legitimate user. This might pose a weakness if facial recognition is used to unlock a smartphone. On the other hand, the failure might not be due to the biometric characteristic itself but to inadequacy of current (sensor) technology.

Behavioral characteristics constitute user actions over time, i.e., for each action, there must be a beginning, an end, and a duration. Consequently, behavioral characteristics indirectly measure properties of the human body. Behavioral characteristics are learned processes and, therefore, can be also re-learned. However, the consensus in the literature seems to be that after reaching a certain age, changes in behavior become more difficult to achieve, even with specific and sustained effort [Woodward et al. 2003]. Behavioral characteristics can therefore be regarded as valid means of identification, even though they are mostly neither as unique nor as permanent as their physiological counterparts. An advantage is that they are less invasive and therefore more user-friendly. For example, a system that analyses keystroke timings or speech patterns can usually do so in the background. In contrast, an iris or fingerprint scan requires specific user actions.

2.2. Requirements and Goals for Pulse-Response Recognition

In this paper, we explore body impedance as a novel biometric characteristic. Body impedance, also referred to as bioimpedance, measures and quantifies the electrical impedance of (parts of) the human body [Martinsen and Grimnes 2011]. We use *pulse-*

response as an instance of a particular body impedance measurement, which is acquired by sending a low-voltage electric signal from the palm of one hand to the other.

When assessing this new biometric mode and envisioning pulse-response based systems it is advantageous to consider lessons learned from past and current biometric systems. General design goals for biometric systems can be found in the literature, e.g., [Jain et al. 2011].

Our requirements and goals we assess the envisaged systems against are described in the following:

Universal. The biometric mode must be universally applicable, to the extent required by the application. The recognition method should apply to everyone intended to use the biometric system.

Unique. The biometric trait must be unique within the target population.

Permanent. The biometric trait must remain consistent over the period of use. Few biometric characteristics stay constant over a lifetime, but they work well if they are consistent over the lifetime of the biometric system,

Unobtrusive. Biometric recognition should be as unobtrusive as possible. If the user can be identified passively, without interference, a biometric system is more likely to be accepted.

Difficult to circumvent. Users of a biometric system should be unable to change the characteristic that is captured for biometric recognition. At a minimum, it should be difficult for a user to modify the biometric characteristic to match that of another user.

Other common non-technical but important goals are:

Acceptability. The biometric recognition should be one that users are likely to feel comfortable with. Clearly, acceptability is a sensible requirement. The capture of pulse-response requires electricity which naturally raises concerns about safety. In Section 3.2 we demonstrate that measuring pulse-response is harmless to health and discuss acceptability and perception of pulse-response based recognition.

Cost effectiveness. The relationship between the distinguishing power of the biometric and its deployment and maintenance costs. Since we focus on assessment of a new biometric mode and are building a prototype, it is premature to seek insights about costs of a possible commercial system.

3. PULSE-RESPONSE RECOGNITION

Pulse-response recognition works by applying a low voltage pulse signal to the palm of one hand and measuring the body's electrical response in the palm of the other hand. The signal travels up through the user's arm, across the torso, and down the other arm. The biometric characteristic is captured by measuring the response in the user's hand. This response is transformed to the frequency domain via the Fast Fourier Transform (FFT). This transformation yields the individual frequency components (bins) of the response signal, which form the biometric features that are then fed to a classifier. Working in the frequency domain eliminates any need for aligning the pulses when they are measured. Details of our measurement setup and experiments can be found in Section 6.

The main reason for the ability of this biometric trait to distinguish between users is due to subtle differences in body impedance, at different frequencies, among different people. When a signal pulse is applied to one palm and measured in the other, the current travels through various types of body tissues – blood vessels, muscle, fat tissue,

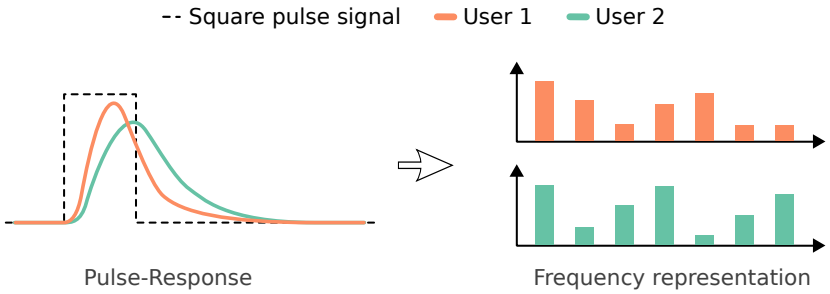

Fig. 1. Overview of pulse-response recognition. The electric response is captured and transformed to the frequency domain. Each individual has a distinct pulse-response due to differences in body impedance.

cartilage and bones – to reach the other hand. Differences in bone structure, muscle density, fat content and layout (and size) of blood vessels result in slight differences in the attenuation of the signal at different frequencies. These differences show up as differences in the magnitude of the frequency bins after the FFT. This is what facilitates distinguishing among individuals. Figure 1 illustrates the concept of how pulse-response recognition works and exemplifies the differences in pulse-response for two different users.

Pulse-response is a physiological characteristic since it measures body impedance which is largely distinct from behavioral aspects. However, it has an attractive property normally associated with behavioral characteristics: it can be captured in a completely passive fashion. Although other physiological characteristics used for biometric recognition also have this feature, e.g., face recognition, pulse-response recognition is not easily circumventable. This combination of unobtrusiveness and difficulty to circumvent makes it an attractive identification mechanism. Essentially, it offers the desirable properties of both physiological and behavioral traits.

At the same time, pulse-response recognition requires special-purpose hardware, which is also true for any other physiological trait. For example, fingerprints need a fingerprint reader, face recognition requires a precision camera and hand geometry – a scanner. Since pulse-response is captured using electrical signals, there are few restrictions on the exact construction of the biometric capture hardware. We explore this issue in Sections 4 and 5.

3.1. Liveness and Replay

A common problem with many biometric systems is presentation attack detection. A fingerprint reader would want to detect whether the purported user's fingerprint was produced by a real finger attached to a human, as opposed to a fingerprint mold. Similarly, a face recognition system would need to make sure that it is not being fooled by a photo or a 3-D artefact. More details and concrete examples are given in Section 9.

In established biometric systems, presentation attacks are usually addressed via some form of active authentication, e.g., a challenge-response mechanism. In a face recognition system a user might be asked to turn his head or look at a particular point during the authentication process. Although this reduces the chance of a photo passing for the real person, the user is forced to take active part in the process, which can be disruptive and annoying if authentication happens on a continuous basis.

In the context of pulse-response recognition, unlike fingerprint or face recognition, it is difficult (yet not impossible) to separate the biometric characteristic from the individual to whom it belongs. If the adversary manages to capture a user's pulse-response on some compromised hardware, successfully presenting it to a sensor would require specialized hardware that mimics the exact impedance of the original user. We

believe that this is feasible: the adversary can devise a contraption that consists of adhesive-covered electrodes attached to each finger-tip (five for each hand going into one terminal) with a single wire connecting the two terminals. The pulse response of the electrode-wire-electrode has to exactly replicate that of the target user. Having attached electrodes to each finger-tip, the adversary can type on the keyboard and the system could thus be effectively fooled. However, the effort required is more than in cases of facial recognition or fingerprints, which are routinely left – and can be lifted from – numerous innocuous locations.

Furthermore, in contrast to face or fingerprint recognition, the pulse-response can be made to depend on the capture platform. Thus, even if the adversary captures pulse-response on one piece of hardware, it would not match the user's measurements on a different capturing device. One way to achieve this is to add a specific (frequency-dependent) resistance to the measurement platform. If the adversary uses its own capture system to measure the user, there is an additional signature which is actually part of the pulse-response reader.

Finally, the real power of the pulse-response recognition is evident when used for continuous authentication (see Section 5), whereby, the person physically uses a secure terminal and constantly touches the keyboard as part of routine work. Biometric verification happens on a continuous basis and thus making it infeasible to use the terminal while at the same time providing false input signals to the authentication system. Of course, the adversary could use thick gloves, thereby escaping detection. However, the biometric system will see input from the keyboard without the expected pulse-response measurement to accompany it, and will lock the session.

3.2. Ethics and User Safety

As mentioned above, the pulse-response is captured by applying a low voltage pulse to one hand of the user and measuring the resulting signal in the other. This involves current flowing through the human body. This process naturally raises questions about user safety and ethics. Clearly, these are important issues that we must address. The issue of safety might be compounded by users having undocumented or undisclosed medical conditions, including implantable medical devices, e.g., pacemakers, that may be adversely affected by applying an external signal to the body.

The amount of current that a particular voltage induces in the human body varies from person to person and depends on external conditions. For example, if a subject's hands are wet, overall conductivity is significantly higher (i.e., resistance is lower) than with dry hands. The same is true if the subject's hands have cuts or broken skin close to where the signal is applied. If resistance is lowered, current strength increases according to Ohm's law. Normal resistance of the human body is between $1,000$ and $5,000\ \Omega$. However, even in very extreme conditions, resistance does not drop bellow $500\ \Omega$. With our current limiting resistor of $10k\Omega$ on the signal generator, the worst case current (with 10V test signal) is $10V/10.5k\Omega = 0.95\ mA$, which is bellow the sensitivity limit. The vast majority of subjects were only exposed to a 1V signal, which translates into the worst case current strength of $0.095\ mA$, less than the current flow induced by touching the terminals of a standard 1.5V battery.

Such a current is on the order of what consumer-grade body-fat scales use. Body-fat scales determine body impedance at predefined frequencies (usually $50\ kHz$) by sending an alternating current of up to $0.1\ mA$ through the body. They then estimate body fat percentage based on the measured impedance and additional information such as body height. Since pulse-response recognition uses similar ampere levels and body-fat scales are intended for daily use, we believe pulse-response based recognition is also safe to use over an extended period of time.

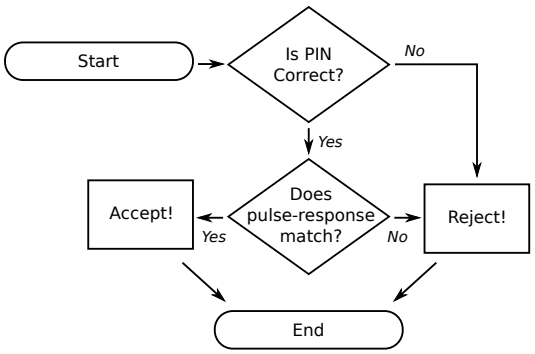

Fig. 2. ATM decision flowchart.

All subjects were given detailed information about the nature of the experiment beforehand and all were given the opportunity to opt out. None expressed any discomfort or, in fact, any perception of the current during the experiments.

Our experimental prototype setup and its safety and methodology have been reviewed and authorized by the Central University Research Ethics Committee of the University of Oxford, under approval reference MSD-IDREC-C1-2014-156.

4. COMBINING PIN ENTRY WITH BIOMETRIC CAPTURE

This section describes the envisaged use of pulse-response recognition to unobtrusively enhance the security of PIN entry systems.

4.1. System and Adversary Models

We use a running example of a metal PIN key-pad with an adjacent metal pad for the user's other hand. The key-pad has the usual digit (0-9) buttons as well as an "enter" button. It also has an embedded sensor that captures the pulse-signal transmitted by the adjacent metal pad. This setup corresponds to a bank ATM or a similar setting.

The adversary's goal is to impersonate an authorized user and withdraw cash. We assume that the adversary can not fool pulse-response recognition with probability higher than that found in our experiments described in Section 7.

We also assume that the ATM is equipped with a modified authentication module which, besides verifying the PIN, captures the pulse-response and determines the likelihood of the measured response corresponding to the user identified by the previously inserted ATM card and the just-entered PIN. This module works as depicted in Figure 2. We assume that the ATM has access to a biometric reference database of valid users, either locally or over a network. Alternatively, the user's ATM card can contain a biometric reference needed to perform pulse-response verification. If stored on the card, this data must be encrypted and authenticated using a key known to the ATM; otherwise, the adversary (who can be assumed to be in possession of the card) could replace it with data matching its own pulse-response.

4.2. PIN Entry Scheme

The ATM has to determine whether a biometric sample acquired from the user while entering the PIN is consistent with the reference in the database. This requires a classifier that yields the likelihood of a sample coming from a known distribution. The likelihood is used to determine whether the newly measured samples are close enough to the reference or template in the database to produce a match. Using our prototype, we can make such decisions with high confidence; see Section 7.

Before discussing security of the pulse-response enhanced PIN entry system, we check whether it meets our requirements stated in Section 2.2.

Universal. A person using the modified PIN entry system must use both hands, one placed on the metal pad and one to enter the pin. This requires the user to actually have two hands. In contrast, a normal PIN entry system can be operated with one hand. Thus, universality of our system is somewhat lower. This is a limitation of the biometric mode, although a remedy could be to store a flag on the user's ATM card indicating that disability, thus exempting this person from the pulse-response verification. This would allow our approach to gracefully degrade to a generic PIN entry system.

Unique and *Permanent*. In Section 7 we show that our prototype can determine, with high probability, whether a subject matches a specific pulse-response. Thus, it is unlikely for two people to exhibit exactly the same pulse-response. We also show that an individual's pulse-response remains fairly consistent over time.

Unobtrusive. In the envisaged setting, the scheme is very unobtrusive, since from the user's perspective, the only thing that changes from current operation is the added requirement to place the free (not used for PIN entry) hand on a metal pad. Naturally, some users might have to change their behavior while operating an ATM, as they could be used to holding something in one hand, e.g., their wallet, or shielding their PIN entry. However, this can be provided for by such a modified ATM. Also, there can be two conductive pads accommodating both left- and right-handed people. In addition, the ATM screen could display system usage instructions, even pictorially to accommodate people who can not read. Similarly, audio instructions could be given for the sake of those who are vision-impaired.

Difficult to circumvent. Given that the pulse-response is unique, the only other way to circumvent it is to provide the sensor (built into the PIN pad) with a signal that would correspond to the legitimate user. Although this is hard to test precisely, assuming that the adversary is unaware of the target user's pulse-response measurements, the task seems difficult, if not impossible.

4.3. Security Analysis of PIN Entry Scheme

The additional layer of security provided by pulse-response recognition is completely independent from security of the PIN entry system alone. Therefore, we model the probability P_{break} that the proposed PIN entry system can be subverted, as:

$$P_{break} = P_{guess} \cdot P_{successful-impostor}$$

where P_{guess} is the probability of the adversary correctly guessing the PIN and $P_{successful-impostor}$ is the average probability that the adversary can fool pulse-response recognition by presenting his own biometric characteristic. In Section 7, we determine the false accept rate to be 9% on average for a zero-effort impostor, i.e., $P_{successful-impostor} = 0.09$.

If a PIN consists of n decimal digits and the adversary has t guesses then $P_{guess} = \frac{t}{10^n}$. Together with $P_{successful-impostor}$ this yields the combined probability:

$$P_{break} = \frac{0.09 \cdot t}{10^n}$$

For example, if the adversary is allowed 3 guesses with a 4-digit PIN, $P_{break} = 2.7 \cdot 10^{-5}$, whereas a 4-digit plain-PIN system has a subversion probability of $3 \cdot 10^{-4}$. Though this improvement might not look very impressive on its own, it is well known that most PIN attacks are performed by "shoulder surfing" or covertly video-taping the PIN entry sequence. These attacks do not involve the adversary guessing the PIN. If we assume that the adversary already knows the PIN, $P_{break} = 9.0\%$ with our system, as opposed to 100% without it.

5. CONTINUOUS AUTHENTICATION

We now present a continuous authentication scheme. Its goal is to verify that the same user who initially (and securely) logged into a secure terminal[1], continues to be physically present at the keyboard. Here, pulse response recognition is no longer used as an additional layer of security at login time. Rather, the user's pulse-response is captured at login time and subsequent measurements are used to authenticate the user by comparing to the initial reference.

5.1. System and Adversary Models

We continue using the example for continuous authentication introduced in Section 1. We use this example throughout this section to make it easier to present the details of the envisaged system. However, applicability of continuous authentication via pulse-response is not limited to this specific scenario.

The system consists of a terminal with a special keyboard that sends out pulse signals and captures the pulse-response. This requires the keyboard to be either made from, or coated by, a conductive material. Alternatively, the pulse signal transmitter could be located in a mouse that the user operates with one hand and the keyboard captures the pulse-response. Without loss of generality, we assume the former option. The keyboard otherwise operates normally and is used for both login and routine activity at the terminal.

We assume that the adversary, with or without consent of the authorized (at login time) user, physically accesses the unattended terminal and attempts to proceed within an already-open session. In security research literature, this attack scenario is also known as "lunchtime" attack (see, e.g., [Eberz et al. 2015]). We assume that the adversary at the keyboard has full access to the active session, and that this happens some time after the original user logged in. Our goal is to detect that the original user is no longer present, and that the keyboard is operated by someone else. If a different user is detected, the system consults a policy database and takes appropriate actions, e.g., locks the session, logs out the original user, raises alarms, or notifies system administrators.

In addition to the peripherals required to capture the pulse-response signal, the continuous authentication system consists of a software process that manages initial login and frequency of periodic reacquisition of the biometric characteristic. This process is also responsible for displaying user warnings and reacting to suspected violations. We refer to it as the *continuous authentication process* (CAP) and assume that neither the legitimate user nor the adversary can disable it.

5.2. Continuous Authentication Scheme

At login time, CAP measures and records an initial pulse-response measurement of the authorized user. Periodically, e.g., every few seconds, CAP reacquires the biometric characteristic by sending and receiving a pulse signal through the keyboard. Each newly acquired measurement is checked against the value acquired at login. If the new measurement is sufficiently distinct from that sampled from the original user, CAP consults its policy database and takes appropriate actions, as discussed above. Figure 3 shows a sample CAP decision flowchart. The decision policy can obviously be further refined. For example, in a corporate setting, all employees could have their biometric template stored in central database to allow for a more thought-out access schema which also includes shared resources or devices.

Before considering the security of the continuous authentication system, we look back at our design goals:

[1]If the measurement apparatus and the electrodes needed to acquire pulse-response readings can be miniaturized, smaller devices such as laptops are imaginable.

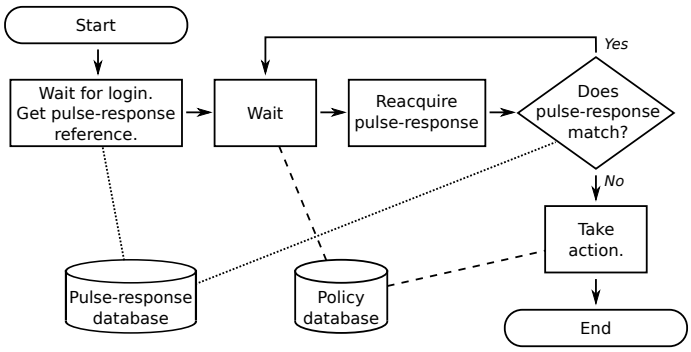

Fig. 3. Flowchart of the *Continuous Authentication Process* decision procedure.

Universal. The users of the system must have two hands in order for the pulse-response biometric to be captured. The same arguments, as in the case of PIN entry, apply here.

Unique and *Permanent*. In Section 7, we show that our prototype can match a pulse-response to previous samples (taken immediately beforehand) with very high accuracy. Average equal error rate is as low as 2%. The fact that the pulse-response reference is taken at the beginning of the session and is used only during that session, makes it easier to overcome consistency issues that can occur when the reference and test samples are days or months apart.

Unobtrusive. Provided the users of the envisioned system periodically come in touch with the electrodes that emit and measure the pulse-response, they do not need to modify their behavior at all and user burden is minimal. In case the electrodes are embedded in a conductive keyboard, this would mean users need to type with both hands. For users who consistently type with only one hand, at least one electrode would have to be incorporated elsewhere, e.g, into the computer mouse the user operates.

Difficult to Circumvent. With a false accept rate of 2% (at equal error rate) it is unlikely that the adversary happens to have a pulse-response similar to the original user and can manage to continuously fool pulse-response recognition. We explore this further in the security analysis section below.

5.3. Security Analysis of Continuous Authentication Scheme

The adversary's goal is to subvert the continuous authentication system by using the secure terminal after the original user has logged in. In the analysis below, we assume that the original user colludes with the adversary. This eliminates any uncertainty that results from the original user "discovering" that the adversary is using its terminal, which is hard to model accurately. We consider the worst-case scenario and the detection probability is a lower bound on security provided by the continuous authentication system. The exact values for the parameters we use in the security analysis are estimated through experiments (based on our data set) that reflect the worst case the proposed scheme could encounter.

We model the security of the continuous authentication scenario with two probabilities. The first is the probability that the adversary is detected immediately, i.e., the very first time when his pulse-response is measured. This corresponds to the complement of the average false accept rate that we report in Section 7 and we call this probability α in the following calculations.

If the adversary's biometric is very close to that of the original user, it might not be detected every time biometric capture is performed. If the adversary manages to fool

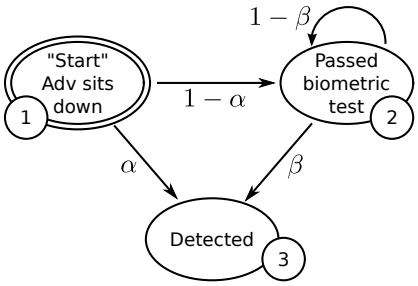

Fig. 4. Markov model of the continuous authentication detection probability. States are numbered 1 to 3 for easy reference in text.

the classifier once, it must be because its biometric characteristic is very close to that of the original user. Thus, the probability that the adversary is detected in subsequent measurements is lower:

$$P[X_i = adv | X_{i-1} = usr] \leq P[X_i = adv]$$

We call this decreased probability β. In Section 8 we will estimate an experimental lower bound for β based on our gathered data set of pulse-response measurements. We measure false acceptance rate in the worst case, i.e., the probability of a successful impersonation attempt for the most promising attacker-victim combination in our data set.

We build a Markov model, shown in Figure 4, with three states to calculate the probability that the adversary is detected after i rounds. When the adversary first accesses the keyboard, it is either detected with probability α or *not* detected, with probability $1 - \alpha$. In the latter case, its pulse-response measurement must be close the original user's. Thus, β is used for the subsequent rounds. In each later round, the adversary is either detected with probability β or *not* detected, with probability $1 - \beta$. To find the combined probability of detection after i rounds, we construct the state transition matrix P of the Markov model, as follows:

$$P = \begin{bmatrix} 0 & 1-\alpha & \alpha \\ 0 & 1-\beta & \beta \\ 0 & 0 & 1 \end{bmatrix}$$

Each row and each column in P corresponds to a state. The entry in row q and column r, p_{qr}, is the probability of transitioning from state q to state r. To find the probabilities of each state we start with a row vector ρ that represents the initial probability of being in state 1, 2 and 3. Clearly, $\rho = [1, 0, 0]$, indicating that we always start in state 1. The probability of being in each state after one round (or one transition) can be represented by the inner product ρP. Probabilities for each subsequent round are determined via another multiplication by P. Therefore, the probabilities of being in each state after i rounds (state transitions), is found as follows:

$$[1, 0, 0] \cdot P^i = [0, (1-\alpha)(1-\beta)^{i-1}, 1 - (1-\alpha)(1-\beta)^{i-1}]$$

As expected, the probability of being in state 1 (the initial state) is 0, since the first state transition forces a transition from the initial state and there is no way back. (See Figure 4.) The probability of being in state 2 (i.e., to escape detection for i rounds) is given by the second element of ρ: $(1-\alpha)(1-\beta)^{i-1}$. The probability of detection is thus: $1 - (1-\alpha)(1-\beta)^{i-1}$. According to our model, using $\alpha = .98$ and $\beta = .36$ (numbers from our experimental results in Section 7 and Section 8) there is a 99.96% chance

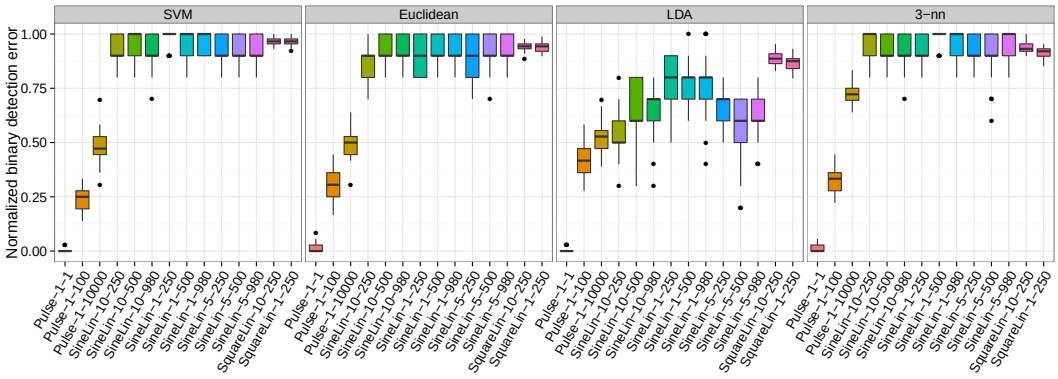

Fig. 5. Box plots of the binary detection error rate for four different classifiers. The distribution shown by each box plot is the result of applying stratified 5-fold cross-validation to the data set five times in a row. We test several different signal types, voltage levels and frequencies for each classifier. We see that narrow pulse signals are consistently performing well.

of detecting the adversary after only 10 rounds. Thus, not surprisingly, acquisition frequency determines the time to detect the adversary.

5.4. Handling False Rejects

False rejects refer to incorrect detection of adversarial presence. If the biometric recognition is used as an additional layer of security during the authentication procedure, this can be managed simply by restarting the login procedure, if the first attempt fails. However, in a continuous authentication setting, where a single (and possibly incorrect) detection might cause the system to lock up, false rejects have to be handled more thoughtfully.

One approach is to specify a policy that allows a certain number of detection events every n-th round, without taking any action. Such a mechanism can help mitigate sensor reading errors or short-term environmental changes that could adversely affect pulse-response recognition and change impedance, such as overly sweaty hands or arms/hands accidentally touching each other or any non-involved metal object.

Another option is to integrate potentially less user-friendly (less transparent) biometric recognition to deal with ambiguous detection events. For example, after a few detection events, the user might be asked to confirm his identity by swiping a thumb on a fingerprint scanner. Such a combined approach would be suitable for a system with very high security requirements. It could employ pulse-response recognition to drastically reduce authentication requests from its principal biometric which might be more obtrusive.

6. BIOMETRIC CAPTURE SYSTEM

In this section, we describe decisions and parameters for our prototype setup that enable us to measure the pulse-response. We conducted several experiments to test different signal types, voltage levels and frequencies.

6.1. Signal Type

Starting out with the hypothesis that body impedance varies depending on frequency and voltage level of the signal, we conducted a preliminary study to test the distinguishing power of various frequency sweeps, and pulse signals with different widths. Although the sweep signals cover a broad range, short square pulse signals prove to be strongly unique among our test subject population.

The box plots in Figure 5 summarize our biometric comparison results with four classifiers that performed well in our application: Support vector machines (SVM), Euclidean distance, linear discriminant analysis (LDA) and 3-nearest neighbors (3-nn). The most promising of the pulses (Pulse), linear sine sweeps (SineLin) and linear square wave sweeps (SquareLin) are listed on the x-axis. The signal name is composed of a signal type, a voltage and a maximum frequency (or width for pulses). The voltage is either 1, 5 or 10 volts. Starting frequency is 1 Hz for all sweeps and the stopping frequency is 250, 500 or 980 Hz, respectively. The width of the pulses (given in hundreds of nanoseconds) is either 100 ns, 10 μs or 1 ms. The y-axis shows the binary detection error rate, i.e., the amount of times the classifier failed to identify a biometric sample correctly, normalized by the number of samples. The distribution denoted by the box plots shows the results of the classifiers achieved by five times 5-fold cross-validation.

We see that the narrow pulse signal outperforms every other signal type by a remarkable margin. We get consistent error rates close to zero for a pulse signal of 1 volt and a width of 100 nanoseconds. Wider pulse signals also give decent results but the quality of the result seems to decrease with the width of the pulse. For the sine and square wave sweeps the results vary significantly with the choice of classifier. Using LDA, some sine sweeps look interesting but nowhere near as good as the narrow pulse signal.

Besides shape and form of the signal, voltage levels are an important factor to consider. It is important that the users of our system do not experience any discomfort when their biometric information is captured. We test three different voltage levels for all signal types: 1, 5 and 10 volts peek-to-peek (Vpp). For sine and square signal sweeps the 10 Vpp and 5 Vpp provides better separation between the subjects but also higher noise levels. For example, in Figure 5, using the LDA classifier, we see that the *SineLin-5-500* signal has a lower detection error rate than the *SineLin-1-500* signal, but the latter has less variance. For pulse signals there is no significant correlation with voltage level. Since the pulse signal is clearly the best choice for our biometric we chose 1 volt pulses to minimize any potential discomfort that users of our biometric system might feel.

6.2. Signal Frequency

We initially thought that (almost) all frequencies would contribute to the distinguishing power of our classifier but our experiments show that the classifier mainly uses the lower frequencies to distinguish between users. In fact, we see an increase in the true positive rate when we only use the first 100 frequency bins of the FFT. This suggests that most of the high frequency content is noise when operating at such low power levels.

6.3. Choice of Classifier

Although we apply an FFT to the data before the classification step we can think of our task as time series classification. This is because an FFT is a reversible linear transformation so the euclidean distance metric is preserved. Thinking of the problem as a time series clustering problem, there are many known approaches that work well. One common method is to compare the first n frequency components by using appropriate distance- or similarity metric. We compare several different classification techniques to see which ones provide the best results for our application.

Euclidean Distance (Euclidean). A new measurement is treated as an n dimensional point and classified according to the euclidean distance to the centroid of each class. This classifier is conceptually very simple but still offers reasonably good results.

Mahalanobis Distance (MH). Rather than assuming uniform and orthogonal dispersion among the frequency components (as in the Euclidean classifier) the covariance matrix for each class is taken into account in the distance calculation. This allows

for a distance metric that is proportional to the shape of the class (in n dimensional feature space). The performance of this classifier improved significantly from the Euclidean distance metric, suggesting that the shape of each class has to be taken into consideration.

Support Vector Machine (SVM). For each pair of groups we train one binary Support Vector Machine classifier (one-against-one approach). The final prediction is found by voting. The inverse kernel width for the Radial Basis kernel is determined by the 0.1 and 0.9 quantile of the pairwise Euclidean distance between the samples. This classifier gives consistently good results and is our final choice of classifier when pulse-response recognition is used for authentication.

Linear Discriminant Analysis (LDA). LDA seeks to reduce the dimensionality of the input data while preserving as much of the class distinguishing power as possible. This classifier turns out to be especially useful for identification. It does however not prove as powerful as the SVM classifier for the binary classification task of authentication.

k *Nearest Neighbor (k-nn).* We test the k nearest neighbors classifier for $k = 1$ and $k = 3$, using euclidean distance. It is a simple classifier that often works very well in practice. In our case though the performance of k-nn is still not as good as SVMs or LDA, respectively.

6.4. Proof-Of-Concept Measurement Setup

In order to gather stable and accurate pulse-response measurements we build a data acquisition platform consisting of: (1) an arbitrary waveform generator, (2) an oscilloscope, (3) a pair of brass electrode handles, and (4) a desktop computer to control the apparatus. Figure 6 is a photo of our setup. We use an Agilent arbitrary waveform generator as the source of the pulse signal. Flexibility of the waveform generator is useful during the initial design phase and allows us to generate the required pulse waveforms in the final classifier. To measure the pulse waveform after the signal passes through a test subject we used an Agilent digital storage oscilloscope which allows storage of the waveform data for later analysis. The output of the waveform generator is connected to a brass handle that the user holds in the left hand. The other brass handle is connected to the oscilloscope signal input terminal. When a test subject holds one electrode in each hand the signal travels from the generator through the body and into the oscilloscope. To ensure exact triggering, the oscilloscope is connected to the synchronization output of the waveform generator.

We use polished brass hand electrodes to ensure optimal electrical contact between the measurement setup and the user. This reduces contact resistance and increases the stability of the measurements.

The function generator and oscilloscope are controlled by a desktop computer that is connected via USB. We wrote a custom software library to set measurement parameters and retrieve the desired waveform data. This software is available upon request.

When measuring the biometric we make each subject follow a specific procedure to ensure that only minimal noise is introduced into the measured data. The test subjects are given a brief explanation of the setup and purpose of the experiment and then told to grab a hold of the brass hand electrodes. The red lead in the left hand and the black in the right hand. The test subjects can choose to either stand or sit in a chair while holding the electrodes as long as they do not touch the sides of their body with their elbows or upper arms. We do this to ensure that the current of the pulse signal has to go through more or less the same path, for all samples and all users. Before each new test subject is measured, the brass handles are wiped down with a disinfectant, both

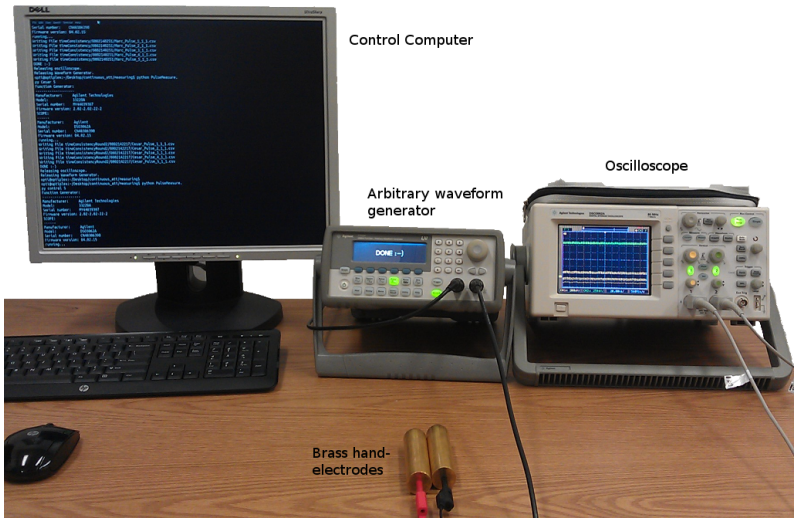

Fig. 6. Proof-of-concept measurement setup. The test subject holds two brass electrode handles [Lyra Nara 2013] and the pulse signal is generated by an Agilent 33220A (20 MHz) arbitrary waveform generator. The receiver is an Agilent DSO3062A (60 MHz), 1 GSa/s digital storage oscilloscope.

for hygienic reasons and to ensure good electrical contact between the electrode and the user's palms.

While our prototype setup ensures accurate biometric measurements and shows feasibility of pulse-response recognition, it might not translate directly to the described application scenarios of PIN entry and continuous authentication in terms of electrode design and other ergonomic requirements. Obviously, further practical tests would be needed before deploying pulse-response recognition, as to find out to what degree soiled electrodes or a change in posture have an effect on the biometric reading.

6.5. Test Subject Population

In the initial design phase, each test subject was sampled ten times for each of the different signal types, for each voltage level and for various frequencies. Once we selected the pulse signal with the best results, samples were acquired for two data sets.

The first consists of 20 samples for each subject, taken in one measuring session. A total of 30 people were measured for this data set, including 9 women and 21 men. We call it the snapshot data set.

The second data set includes 25 samples per subject from a total of 16 subjects, obtained in five different sessions over time. To assess stability and permanence of pulse-response, we measured the biometric over a longer period of time. We sampled all test subjects at different times during the day over the course of several weeks. The median timespan between consecutive sessions was 8 days and there was a minimum time interval of at least one day between sessions.

We tried to sample subjects in order to end up with sampling conditions as diverse as possible, for each subject, to capture various other potential factors that might influence body impedance, such as varying body water percentage, body temperature or time of the day.

Table I summarizes the composition of the test subject population.

Table I. Test subject population and sample size

Data set	Test subjects	Females	Males	Samples per subject
Snapshot	30	9	21	20
Over-time†	16	2	14	25

The age band of the subject population ranges from 24 to 38.
†Test subjects were measured in five different sessions over time.

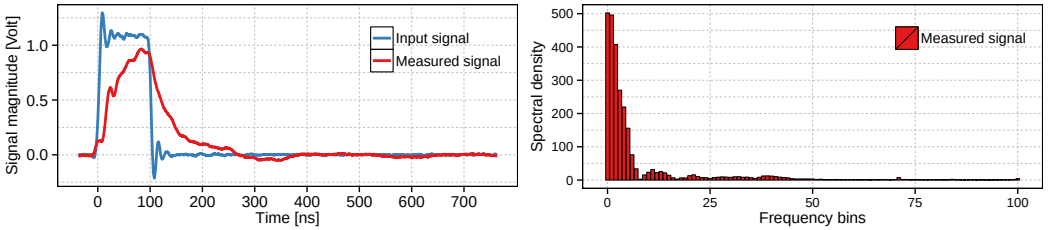

(a) Input and output waveforms

(b) Raw FFT data of the measured pulse

Fig. 7. One measurement consists of 4,000 samples with the rate of 500 MSa/s. In Figure 7a it is apparent that the measured pulse has been modified by passing through the user. The FFT data shown in Figure 7b consists of the first 100 frequency bins of the measured waveform.

6.6. Feature Extraction

Data extracted from the measurement setup is in the form of a $4,000$ sample time-series describing voltage variation as seen by the oscilloscope. Figure 7a shows the input pulse sent by the waveform generator and the pulse measured by the oscilloscope. Time series measurements are converted to the frequency domain using the FFT and the first 100 frequency bins of the FFT data are used for classification. Operating in the frequency domain has several advantages. First, there is no need to worry about alignment of the measured data pulses when computing metrics, such as the euclidean distance between pulses. Second, it quickly became apparent that only lower frequency bins carry any distinguishing power. Higher frequency bins were mainly noise, meaning that the FFT can be used to perform dimensionality reduction of the original $4,000$ sample time-series to the vector of 100 FFT bins. Figure 7b shows an example of the raw data we end up with after the FFT. This data is then fed into the classifier.

6.7. Classification Performance Metrics

We use false accept rate (FAR) and the false reject rate (FRR) as binary classification performance metrics to assess system performance of our prototype setup.

To illustrate the FAR and FRR graphically we draw the receiver operating characteristic (ROC curve) which shows the relationship between these two performance numbers. The ROC curves shown in the following are vertical averages. We compute a ROC curve for every test person and calculate the average over all false reject rates for given false accept rates (see [Fawcett 2006] for an algorithm on vertical aggregation of ROC curves).

A common performance metric for biometrics is the equal error rate (EER). It denotes the rate at which errors for acceptance and rejection are equal and is a straightforward way to compare different ROC curves. Equal error rates for the best performing classifiers will be presented in Section 7.

To assess the performance of pulse-response recognition in identification, we compute the ranking success $\mathrm{Rank}(N)$. The ranking success is a metric that measures the ratio of query samples for which the corresponding template is amongst the first N templates

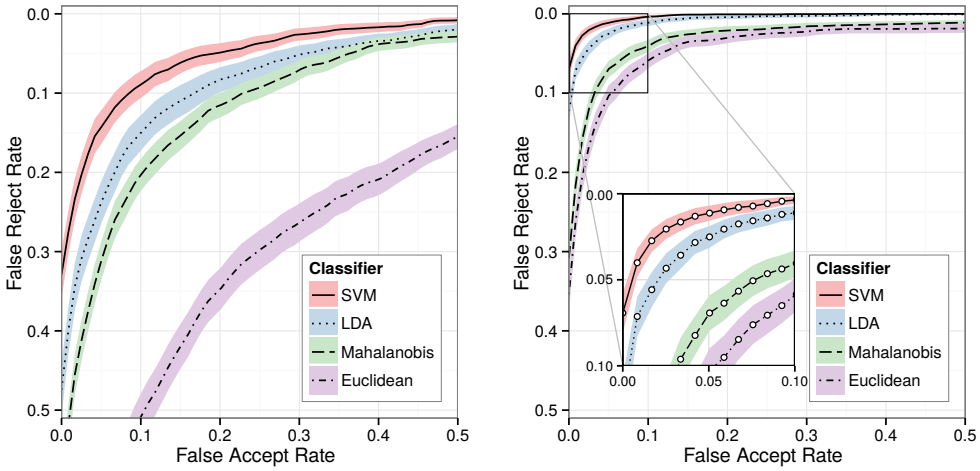

(a) Over-time data set (b) Snapshot data set

Fig. 8. Receiver operating characteristic for authentication. The results presented are averages over all users and obtained by applying 5 times stratified 5-fold cross-validation (ROC curves are vertical averages). Shaded areas show the 95% confidence interval for each classifier.

out of all stored templates in the database if the templates are sorted in decreasing order according to their similarity values. Ideally, $\mathrm{Rank}(1) = 1.0$, which means that for all query samples the corresponding template from the database has been assigned the highest similarity value.

To obtain unbiased and realistic performance measurement numbers, the data sets are partitioned into learning set and test set. We make sure that the test set for the over-time data spans all five measurement sessions. For both data sets, the partitioning into training and test set is repeated multiple times by stratified cross-validation to acquire a robust estimate of the performance of the biometric modality.

7. EXPERIMENTAL RESULTS

In this section we present the results of our experiments with pulse-response recognition, a narrow pulse signal, that resulted from our analysis as the final biometric characteristic. The design decisions and motivations behind selecting a short square pulse signal are described in detail in the previous section. We report system performance figures of various classifiers when they are applied to pulse-response recognition. To be precise, we divide the results into two different types of classifiers according to the usage scenario of the biometric trait. We present classifiers for authentication and for identification.

We sub-divide the results into the two underlying data sets: (1) those from the snapshot data set, which show the inherent distinguishing power of the pulse-response, and (2) those based on the data sampled over time, which assess stability (permanence) of the pulse-response.

Within our data set and due to our straightforward feature extraction, we did not experience any failure to enroll or failure to capture errors, which means the classifier performance corresponds to the actual system performance of our prototype setup (FAR and FRR). We therefore do not report classifier and system performance separately.

7.1. Authentication Classifier

Authentication is a binary classification task. The classifier has to decide whether a presented sample belongs to the group of samples reflecting a specific user or not. An authentication classifier for pulse-response recognition is used in the running example of Section 4 where pin entry is combined with biometric measurements and in Section 5 where pulse-response measurements enable continuous authentication. To simulate an authentication procedure with pulse-response recognition, we separate the samples into two classes: Samples belonging to the legitimate user and samples from all other users. Samples from other users are collected in a large pool and represent potential impostors. Once the classifier is trained, it is presented with unseen samples from both classes. Then FAR and FRR are computed on the basis of the classifiers' prediction.

In order to solve the binary classification problem we test four of the classification algorithms described in Section 6.3: Support Vector Machines (SVM), Linear Discriminant Analysis (LDA), Mahalanobis distance (Mahalanobis) and Euclidean distance (Euclidean). Figures 8a and 8b show the performance of each of these methods when applied to the over-time data set and the snapshot data set. The depicted ROC curves are averages over all test subjects and describe the relationship between the FRR on the y-axis and the FAR on the x-axis. If a higher FRR is acceptable, a lower FAR can be achieved and vice versa, i.e., if a lower FRR is required, the classifiers show a higher FAR. By changing the discrimination threshold the classifiers can operate on any point on the curve if desired.

To ensure statistical robustness the ROC curves are constructed by performing 5-fold cross-validation and averaging the results vertically. The confidence intervals reveal that there is very little variance in the classifiers' performance even if the data set is partitioned into different training and test sets.

The ROC curves show that all subjects are recognized with high probability, as the FRR and the corresponding 95% confidence intervals confirm. SVM outperforms all other classification techniques, followed by LDA and Mahalanobis. SVM achieves a FRR of less than 10% and a FAR of less than 10% at the same time, i.e., an EER of 10%. Given that this assessment is based on the over-time data set it is a remarkable result.

Applying the classifiers to the snapshot data set yields even better performance as Figure 8b reveals. At a FAR of 5%, FRR is close to 100% when using SVM as classifier. This result suggests that pulse-response is a very viable biometric characteristic for continuous authentication and shows remarkable distinguishing power. In a continuous authentication system where a certain percentage of false rejects (incorrect rejection of a legitimate user) can be accepted — such as the one described in Section 5 — pulse-response recognition will, with high probability, detect all adversarial samples.

Moreover, pulse-response recognition seems to be especially effective as a biometric trait if the stored biometric template is fairly recent in relation to the measurements it has to identify. All classifiers show a significant improvement in performance if they only have to deal with samples from a single measurement session, i.e., the snapshot data set. Performance on the over-time data set is likely to be improved with more measurement sessions. The classifiers will gain a clearer picture of the variability of each subject's body impedance if they have access to samples from additional points in time.

7.2. Identification Classifier

Biometric identification is a multi-class classification problem. The goal is to identify a person as accurately as possible given unlabeled biometric samples.

We test five different classifiers in the identification scenario. The conceptual generalization to the multi-class setting is straightforward for all classifiers: The Euclidean

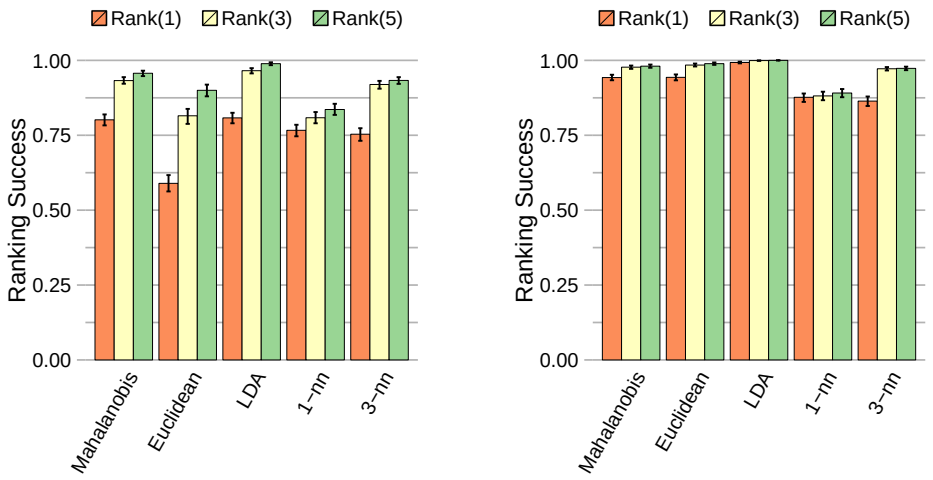

(a) Over-time data set (b) Snapshot data set

Fig. 9. Ranking success rates for identification. The results presented are averages over all users and obtained by applying 5 times stratified 5-fold cross-validation. Error bars show the 95% confidence interval for each classifier.

and Mahalanobis distance classifiers increase the number of centroids to one centroid per class. LDA generalizes to Multiclass-LDA by introducing one mean per class and measuring between-class variability through the covariance matrix of the class means.

Ranking success rates obtained from the identification classifiers are shown in Figures 9a and 9b. We depict $\mathrm{Rank}(1)$, $\mathrm{Rank}(3)$ and $\mathrm{Rank}(5)$. The classifiers have been trained on both, the over-time data set (Fig. 9a) and the snapshot data set (Fig. 9b). The ranking success rates illustrated in the bar plots are averaged over all subjects and obtained by applying 5-fold cross-validation, similar to the authentication scenario.

Even with the increased complexity of multiple classes, all tested identification classifiers perform reasonably well on the over-time data and very well on the snapshot data set. For the snapshot data set a ranking success close to $\mathrm{Rank}(1) = 1.0\%$ is possible using Multiclass-LDA as classifier. The nearest neighbor classifiers (1-nn and 3-nn) can not quite reach the performance of the Mahalanobis distance method and Multiclass-LDA. Clearly the conceptually simpler Euclidean distance method can not cope with the added variability present in the over-time data set (see Fig. 9a).

All classification methods benefit from measurements that are acquired in a relatively short time frame, e.g., the snapshot data set. They can improve performance significantly if trained and tested on these samples only. Measurements taken far apart are influenced by very different conditions. There might be physiological changes, such as weight loss or gain, or there might be differences in the ambient temperature, humidity, clothing, or a number of other factors. The added uncertainty becomes apparent in the classification performance and in turn effects the ranking success rates (compare Fig. 9a with Fig. 9b).

7.3. Summary

Table II summarizes the results of the best classifiers for authentication and identification achieved with our prototype setup, on both, the snapshot data set and the data set taken over time. For authentication SVM gives the best results whereas for identification Multiclass-LDA proves to be the most suitable classifier.

Table II. System performance of the prototype setup averaged over all users in [%]

Authentication (SVM classifier)	FAR	FRR	Accuracy	EER
Snapshot set	2	2	96	2
Over time	9	9	87	9
Identification (LDA classifier)		$\mathrm{Rank}(1)$	$\mathrm{Rank}(2)$	$\mathrm{Rank}(5)$
Snapshot set		99	100	100
Over time		81	97	99

Performance metrics are calculated using five times 5-fold stratified cross-validation. Values shown reflect the performance achieved with the best classifier for each scenario.

In an authentication scenario pulse-response recognition achieves a very low EER of 2% on the snapshot data set and an EER of 9% on the over-time data. This makes it clear that pulse-response is a viable biometric characteristic for authentication.

If pulse-response recognition is used for identification purposes a ranking success of almost 100% can be achieved for the static snapshot data set. According to our experiments, even if the biometric measurements are captured in sessions that are weeks apart pulse-response recognition will reach a ranking success rate $\mathrm{Rank}(1)$ of 81% and 97% for $\mathrm{Rank}(3)$, respectively.

8. IMPERSONATION OF PULSE-RESPONSE

In this section, we introduce an impersonation attack and measure similarity of pulse-response samples. We experimentally estimate worst-case probabilities for different scenarios where an attacker could impersonate a legitimate user by fooling the biometric system using his own biometric measurements (zero-effort impersonation).

8.1. Attacker Model

We consider four attack scenarios relevant to pulse-response recognition. Similar to the previous section we differentiate between authentication and identification. In addition, a potential attacker who tries to impersonate a legitimate user may or may not be known to the system. We refer to an attacker whose biometric template is known as an *internal* attacker and if no biometric template or reference is know, we call it an *external* attacker. Thus, an internal attacker has been registered and is enrolled in the system. An external attacker has never used the system and no pulse-response measurements have been gathered.

Regardless of its type, the attacker's goal is to impersonate a legitimate user of the system. The attacker tries to achieve this by using its own pulse-response and trick the classifier.

To give a realistic experimental lower bound on the attack probabilities for zero-effort impersonation, we base our analysis on the over-time data set. The results in Section 7 made evident that classifying pulse-response samples with increased variability is more challenging. Consequently, we assume that it is also more difficult to detect an attacker under these conditions.

The attack scenarios are limited by the scope of our data set but they nevertheless provide an accurate view on the behavior of pulse-response recognition.

8.2. Internal attackers

We take our best classifiers from Section 7 and estimate their performance for all possible attacker-victim pairs in our test subject population. We first train the classifiers on the entire data set and then ask them to classify biometric samples from a predefined

Table III. Average and worst-case performance for attack scenarios

	internal attacker [%]		external attacker [%]	
	average case	worst case	average case	worst case
Authentication (SVM classifier)				
– Sensitivity	98.8	76.0	98.8	88.0
– Specificity	99.9	99.0	95.0	36.0
Identification (LDA classifier)				
– Sensitivity	99.9	76.0	99.6	80.1
– Specificity	99.5	99.0	99.0	92.0

Average and worst-case sensitivity and specificity for four attacks scenarios. We distinguish between authentication and identification and between internal and external attackers. External attackers are not known to the classifiers.

attacker and a predefined victim only. We thereby measure sensitivity (that corresponds to the complement of FAR, i.e., $1 - FRR$) and specificity (which denotes the complement of FAR, i.e., $1 - FAR$) for a specific combination of attacker and victim.

This performance assessment is repeated for all possible attacker-victim combinations which lets us compute average as well as worst-case performance of sensitivity and specificity. The results can be found in Table III in the column labeled *internal attacker*. Not surprisingly, average sensitivity and specificity attain very high numbers and confirm our previous findings about the classification power of pulse-response recognition. Consistent specificity of almost 100% – on average and in the worst case – guarantees that an internal attacker is very likely to be detected, whether pulse-response recognition is used for authentication or identification.

Sensitivity seems to vary slightly more than specificity. For certain attacker-victim combinations sensitivity only reaches 76%. This means that a particular legitimate user is recognized in 76% of the tested samples. In all remaining cases he was rejected because the classifier mistook him for the attacker. These numbers are congruent with our results from Section 7 where we discover that more variability in the pulse-response measurements affects average sensitivity to a greater extent than average specificity.

8.3. External attackers for authentication

To model an external attack on pulse-response recognition we pursue a similar procedure as outlined above for internal attackers. The main difference is that no attacker samples are included in the training phase of the classifiers. The classifier should be able to identify adversarial samples without knowing a template describing the attacker's pulse-response measurements.

We exploit the nature of the binary classification problem of authentication and form two classes of samples. Having set aside all measurements from the attacker, we define a class containing samples from the victim and a second class consisting of samples of any other user. Although the actual attacker is not represented in this pool of training samples the classifier can gain a good understanding of what measurements other than those from the victim look like. During classification, an external attacker is likely to fall into the group of "other" users despite the fact that the classifier does not see any of the attacker's samples during training. In Figure 10 we show a graphical representation of all possible attacker-victim combinations for an external attack on the authentication classifier. Sensitivity and specificity are shown for each attacker-victim pair, as well as average sensitivity and average specificity. The resulting matrix appears to be nearly symmetric. If two subjects have similar pulse-response measurements it is almost equally likely for both of them to be able to successfully impersonate the other. There are a few deviations, however, which for instance include subjects *Remo* and *Mason* (see Figure 10). *Remo* has a higher chance of impersonating *Mason* than the other

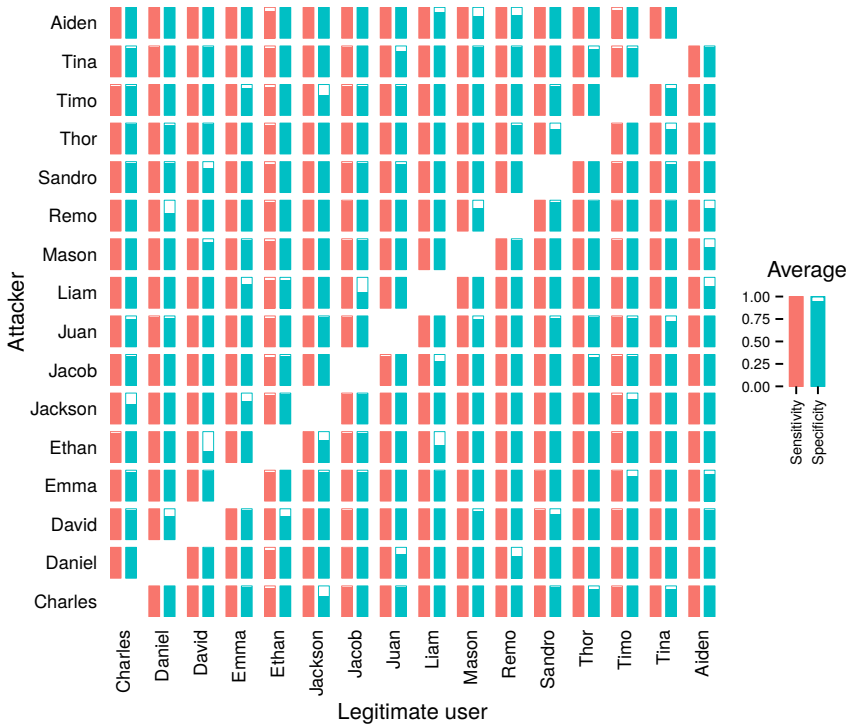

Fig. 10. An external attacker tries to impersonate a legitimate user. Sensitivity and specificity for every possible attacker-victim combination of the over-time data set based on unseen samples from both, attacker and legitimate user (test persons have been anonymized with pseudonyms).

way round, as specificity is lower when *Remo* simulates the attacker. These differences stem from the fact that the class of samples from two different users can have different shape and dispersion in the feature space. The classifier will not necessarily create symmetrical decision boundaries when it is trained on different subsets of the data.

From the results in Table III we see that on average the authentication classifier performs almost equally well in both attacker scenarios, internal and external. Sensitivity and specificity are above 95% in all cases. Although average performance is very high, a few attacker-victim combinations reveal detection probabilities significantly below average. For instance, if subject *Ethan* wants to impersonate *David* then specificity is estimated at 36% which will result in a 64% chance for *Ethan* to go undetected and successfully fool pulse-response recognition (see Figure 10). *Ethan* and *David* must have a very similar pulse-response. The fact that some attacker-victim pairs have similar measurements is what motivated the Markov Model in Section 5.3. The model takes into consideration that the measurements of the attacker might be statistically similar to the legitimate user and as a consequence the attacker successfully passes the biometric test at first and only gets caught eventually.

8.4. External attackers for identification

When pulse-response recognition is used for identification, reliable detection of external attackers becomes more intricate. The classifier has to distinguish between multiple classes and detect attacker samples at the same time. It is possible to construct a binary classifier for every single user which decides between legitimate user and attacker. However, this approach requires an aggregation scheme that collects the classifiers'

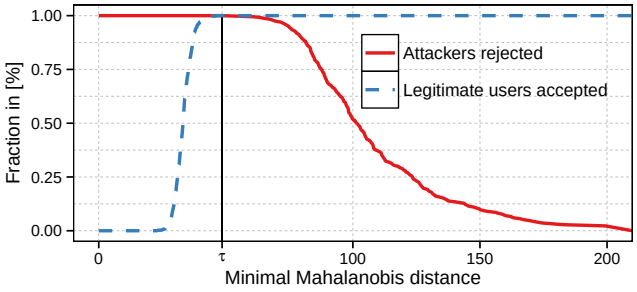

Fig. 11. Detecting external attackers as statistical outliers with minimal Mahalanobis distance between sample and class means. At discrimination threshold τ the fraction of detected attackers is equal to the fraction of recognized legitimate users.

outputs to produce the final decision whether the presented measurement is indeed an attacker or not. Results in Section 7.2 showed that for pulse-response recognition pairwise SVM classifiers do not perform as well as conceptually simpler methods, such as Multiclass-LDA. Starting out with this insight we opt for a less complex model to detect attacker samples. It is based on the assumption that samples from unknown subjects can be detected as statistical outliers. Samples from an external attacker originate from an unknown source as no such samples have been seen by the system before classification. These adversarial measurements might not share any statistical characteristics with the measurements the classifier has encountered during training phase.

An effective approach to an outlier detection scheme is to determine mean and covariance for each class of samples representing a registered user. This information is used to compute the Mahalanobis distance between a new measurement and all stored biometric templates, i.e. the Mahalanobis distance to the class means. Should a new measurement happen to be far from all class means then the likelihood of it being an attacker sample is high. If the minimal distance exceeds a certain threshold the sample is declared as an attack and filtered out.

The described method is essentially the same as the Mahalanobis classifier that we tested for pulse-response recognition in Section 7.2. This time though, we do not assign class labels to the samples but rather compute the likelihood (i.e., the distance) that a sample belongs to any of the stored templates. The motivation behind choosing Mahalanobis distance for outlier detection is twofold: It showed very good classification results for pulse-response measurements and its application to outlier detection is straightforward. There is no need to train multiple classifiers and no additional class to accommodate outliers is required.

The performance graph in Figure 11 shows the discrimination threshold for the minimal Mahalanobis distance against the result of the outlier detection. Varying the discrimination threshold not only has an effect on how many legitimate users are recognized but also on how many attackers are detected. Ideally, the system would reject all attacker samples and accept all samples from registered users. The threshold $\tau = 48.5$ used in the experiments is chosen in such a way that the percentage of detected attackers is equal to the percentage of correctly identified legitimate users. τ is found by 2-fold cross-validation and achieves an error rate of almost 0%.

After having filtered out the attacker samples the system continues to assign class labels to the remaining measurements – those which have not been found to represent an external attacker. The classification of these samples is analogous to the scenario where only internal attackers are considered. The biometric system employs a Multiclass-LDA

classifier to solve the classification task, similar to the identification classifier found in Section 7.2.

Since the system now contains two sources of possible misclassification errors (the classifier for the user samples and the preceding outlier filtering) the performance assessment must make sure to take this fact into account. In particular, we need to consider the rejection of legitimate users during outlier detection. A legitimate user who is incorrectly identified as an external attacker must be treated as a wrong assignment and should impact the sensitivity score.

Table III lists sensitivity and specificity metrics the identification classifier is able to achieve when samples are pre-filtered by Mahalanobis distance to detect external attackers. Performance experiences almost no decrease compared to the scenario for internal attackers. Average sensitivity and average specificity stay at a very high level of 99%. Worst-case sensitivity is even increased from 76.1% to 80.1%. Worst-case specificity is affected to a marginal extent. It changes from 99% to 92% which supports our initial assumption that it is more challenging to detect external attackers than internal attackers. We can still conclude, however, that there is a high chance that impersonation attempts from external attackers are detected. This is mainly due to the effective outlier detection scheme which filters out attacker samples before the measurements are fed to the classifier.

9. RELATED WORK

Biometric characteristics, as a means of recognizing an individual using physiological or behavioral traits, has been an active research area for many years. A comprehensive survey of established physiological biometrics can be found in [Jain et al. 2006]).

While physiological biometrics tend to be relatively stable over time, they can be sensitive to deception and presentation attacks. These include, for instance, attacks on: (1) fingerprint identification, e.g., using mock fingers made of glycerin, gelatin or silicon [Barral and Tria 2009; VIRDI Biometric 2009], (2) facial recognition, e.g., using photographs or 3D models of an actual user [Nguyen and Bui 2009; Boehm et al. 2013], and (3) iris scan, e.g., using patterned contact lenses that replicate a genuine user's iris [Galbally et al. 2012].

In contrast, behavioral biometrics are thought to be harder to circumvent. However, the performance of systems that implement behavioral biometrics, in terms of false rejection rates (FRR) and false acceptance rates (FAR), is usually lower and can require re-calibration due to varying and often erratic nature of human behavior. Initial results on behavioral biometrics were focused on typing and mouse movements, see, e.g., [Spillane 1975; Clarke and Furnell 2007].

In particular, keystroke dynamics gained lots of popularity through [Monrose et al. 1999], where it was used to augment password authentication similarly to our PIN entry scenario. Keystroke dynamics make use of the typing cadence and timings of an individual while typing on a keyboard and are a biometric recognition method that could be added to our PIN entry scenario as an additional modality. However, as recognition rates of keystroke dynamics greatly improve with longer sampling durations, it would be even better suited to continuous authentication. Keystroke dynamics could serve as an alternative or as a complement to our pulse-response based recognition. We compare keystroke dynamics and pulse-response recognition in Section 9.1.

In contrast to keystroke dynamics, some research studies on mouse movement biometrics argue that it should not be used as biometric for authentication, as it has too high intra-class variability, is highly device-dependent [Pusara and Brodley 2004] and requires a long sampling duration, while others report high accuracies [Nakkabi et al. 2010; Gamboa and Fred 2004; Zheng et al. 2011]. The authors of [Zheng et al. 2011] achieved equal error rates (EER) as low as 1.3% using successive mouse actions

between clicks. Some of the best results has been reported in [Nakkabi et al. 2010] with a FAR of 0.36% and a FRR of 0%, although it has been suspected that this result was influenced by recording the data on a different computer for each user [Jorgensen and Yu 2011].

An evaluation of keystroke dynamics, mouse movements, application usage and system footprint can be found in [Deutschmann et al. 2013]. A total of 99 users participated in the study and biometric data covering 20 hours per week during a span of 10 weeks was gathered. In addition to biometric traits, the system acquired CPU and RAM usage and the computer programs that were used most often. The study comes to the conclusion that keystroke dynamics prove most useful for continuous authentication.

9.1. Comparison to keystroke dynamics

Keystroke dynamics is one of the most researched behavioral biometrics. Some of the first scientific studies that propose to harness the distinguishing capabilities of keyboard characteristics for identity verification date back to the mid 1970s and can be found in, e.g., [Umphress and Williams 1985] and [Spillane 1975]. Since then, many different recognition methods have been proposed. The most straight-forward methods are based on relatively simple statistics, such as mean typing times and their standard deviations [Joyce and Gupta 1990; Araujo et al. 2005]. Over the last few years, several pattern-recognition methods have come into vogue and been applied to keystroke dynamics, such as e.g., neural networks [Cho et al. 2000], fuzzy logic [Tran et al. 2007], and support-vector machines [Giot et al. 2009]. A survey on the large body of literature on biometrics using keystroke dynamics is given in [Joyce and Gupta 1990; Banerjee and Woodard 2012; Teh et al. 2013] and in the comprehensive background section in [Killourhy 2012].

As mentioned, if the keyboard users are typing on is conductive, a biometric system could be designed as to measure both biometric traits, keystroke timings and the pulse-response, at the same time and only with minimal user intervention. Both biometric modes do not require the user to change his normal work-flow when typing on a keyboard which makes the biometric recognition process very unobtrusive. Clearly, these two modalities could complement each other and result in a more powerful biometric system. Unfortunately though, both these modalities have the drawback of not being able to acquire biometric measurements during periods when there is no user input. Assuming users of such a combined system do not rest their hands or fingers on the keyboard while inactive, neither keystroke dynamics nor pulse-response recognition can bridge the breaks between typing phases. In such cases, other recognition methods, e.g., a video camera for face recognition, could be a better complement and increase security to a greater extent than keystroke dynamics and pulse-response recognition in combination with each other. We therefore compare the performance of keystroke dynamics and pulse-response recognition in more detail.

To this end, we evaluate a scenario specifically designed for this comparison. This allows us to compare pulse-response with performance numbers for keystroke dynamics found in literature. We assume that users type on a conductive keyboard and every keystroke results in only one captured pulse-response measurement. Since the square pulse used for the capture has a duration of 100 nanoseconds, many more measurements would in theory be possible during a single keystroke. The enrollment data for this analysis is comprised of five random measurements per user, taken from our over-time data set. The validation data consists of 17 measurements per user, randomly sampled from the snapshot data set. Choosing training and validation data in such a way, we simulate the verification of new measurements (captured in quick succession while the user is typing) with the help of a stored biometric template obtained

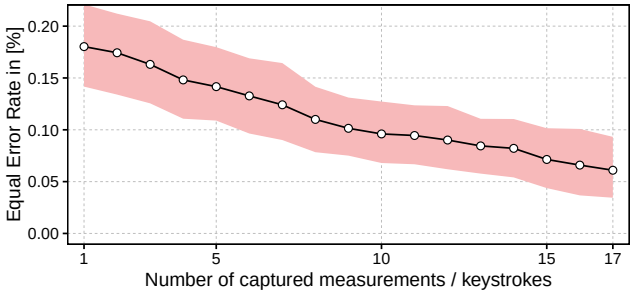

Fig. 12. Equal error rate (EER) of pulse-response recognition in relation to the number of keystrokes. We assume users are typing on a conductive keyboard and every keystroke results in one pulse-response measurement, e.g., for a five-letter word, five measurements can be captured (measuring errors are omitted). The solid line represents mean EER over all users, the shaded area shows the 95% confidence interval.

during enrollment[2]. The final authentication decision is made based on the aggregated classification outcomes of each individual measurement.

Figure 12 shows the equal error rate depending on the number of captured measurements, averaged over all users. The 95% confidence interval of the mean is depicted as a shaded area. We estimated it by resampling the subsets for enrollment and validation data 25 times for each user. After one single measurement, i.e., after one keystroke, mean equal error rate averages to 18.0% and steadily declines to 6.14% if up to 17 subsequent measurements can be captured. The performance of keystroke authentication systems varies in a similar fashion: If verification consists of a single word, i.e., as it is the case in password augmentation, only a small amount of keystroke data can be captured by the system and recognition rates are consequently lower. The study in [Giot et al. 2009] which uses a 16 character pass-phrase for both enrollment and verification achieves an equal error rate of 6.96% (vs. 6.14% of pulse-response recognition) whereas free text recognition (users are allowed to type anything for enrollment and verification) can achieve equal error rates as low as 0.95% [Gunetti and Picardi 2005]. Short typing sequences or passwords, however, yield similar results to pulse-recognition. The study in [Bleha et al. 1990] uses passwords between 11 and 17 characters and resulted in 8.1% FRR and 2.8% FAR. [Araujo et al. 2005] operates with a text length of 10 and achieves a FRR of 11.57% and FAR of 1.89%. Finally, the authors of [Hocquet et al. 2005] are able to get 6.0% FRR and 0.5% FAR while using a text length of 25.

Although research literature has shown that typing patterns between individuals can have similar characteristics, and error rates are low, misidentification is possible as in traditional fingerprinting. Recognition rates are high enough such that keystroke dynamics can be considered unique to each individual [Killourhy 2012; Araujo et al. 2005]. However, there are research studies that question the uniqueness property of keystroke biometrics. The most prominent one is [Tey et al. 2013] where attackers are shown the typing pattern of their victims and make a conscious attempt to imitate. The attackers receive training through a textual and graphical feedback interface. After training, false acceptance rate increases from 0.20 to 0.42 if attackers have partial knowledge of the typing statistics of the victim, and from 0.24 to 0.6 if entire typing statistics are known. These results show that keystroke dynamics might be questionable in high-security environments and existing commercial solutions using keystroke biometrics might not withstand targeted attacks.

[2]This is different from the continuous authentication setting in Section 5 because new measurements are not compared to an initial reference measurement (temporary template) obtained at login time, but validated against a pre-existing, stored biometric template.

9.2. Touch(-screen) biometrics

Nowadays, many modern personal electronic devices, such as smart phones and tablets, usually possess a capacitive touchscreen as input device as opposed to keyboard and mouse. Quite recently research has turned to how to make continuous user authentication work with input signals received from a touchscreen. Touchscreen biometrics, i.e., taps, strokes, swipes and gestures executed by one or multiple fingers on a touchscreen, are similar to keystroke dynamics as they can only be measured and evaluated during active user input. They are considered very unobtrusive as they measure users' touchscreen actions which are part of the natural work-flow when interacting with a smart phone or similar device. If the biometric capture mechanism needed to measure body impedance can be miniaturized in the future, pulse-response recognition might also be used on smaller devices where it could complement touch(-screen) biometrics, similarly to keystroke dynamics.

The first work that thoroughly investigates the applicability of touchscreen input as a behavioral biometric can be found in [Frank et al. 2013]. The authors propose 30 behavioral features that can be extracted from a user's interaction with a smart phone equipped with a touchscreen. The paper concludes that touchscreen features might not be applicable to long-term authentication, they could, however, still serve as part of a multi-modal biometric recognition scheme or secure short absences of usage without immediately locking the device. In [De Luca et al. 2012], for instance, touch characteristics are used to unlock a smart phone and to enhance swipe/shape password patterns for instant authentication. The authors achieved a recognition rate of 57% in a two-day user study.

In [Feng et al. 2012] another framework and a prototype for continuous user authentication on mobile devices is presented. It consists of a sensor glove that delivers fine-grained features, e.g., orientation, direction, rotation, of the finger movements and a smart phone that collects touch gesture data. This augmented approach achieves slightly worse recognition rates, but the authors believe that their system could be used successfully for post-authentication security for a certain amount of time after the user authenticates by some other means, i.e., password or other biometric.

A similar approach is presented in [Holz and Knaust 2015] where a watch-like prototype measures the user's skin impedance profile of the wrist in order to modulate a user-specific signal onto the user's skin that can be picked up by a touchscreen. This allows seamless and transparent authentication on each touch the user makes. The authors recruited 10 participants for a lab evaluation and claim that their classifier produces no false positives when identifying users.

9.3. Body impedance / bioimpedance-based biometrics

[Revett and de Magalhães 2010] covers recent papers on cognitive biometrics based on the electroencephalography (EEG), the electrocardiogram (ECG), and the skin conductance, also called electro-dermal response (EDR), and describes how these biometrics can be harnessed for user authentication. Skin conductance (EDR) is directly related to body impedance in terms of modality and acquisition method. The main difference is that, unlike body impedance, it captures the emotional state of an individual and not necessarily a physiological trait. The resistance of the skin can vary significantly due to the embedded sweat glands which are controlled by the nervous system. Body impedance-based biometric recognition methods (such as pulse-response recognition), on the other hand, focus on extracting physiological characteristics independent of emotional state by measuring entire parts of the human body, not only skin conductance.

Probably the most related to this paper is the work in [Cornelius et al. 2012] where bioimpedance is used as a physiological characteristic. A wearable sensor is designed

to passively recognize wearers based on a body's unique response to the alternating current of different frequencies. The authors design a prototype wristband that captures electrical impedance around at the wearer's wrist, as opposed to measuring body impedance from one hand to the other. Experiments in [Cornelius et al. 2012] were conducted in a family-sized setting of 2 to 5 subjects, where a person wears the bioimpedance sensor on the wrist. They achieve recognition rate of 90%. In a more recent study [Cornelius et al. 2014] the authors improved their prototype and increased the number of test subjects. They report FAR and FRR of 2% for samples taken within a day. Our biometric recognition method solves a different problem—we propose a recognition method that works by temporarily touching two electrodes, not a wearable device—but our technique also uses the body's response to a signal. We achieve a similar error rates when samples are taken in one session and slightly higher error rates when samples are taken weeks apart.

Although not directly related to our work, it is interesting to mention a cryptographic key generation scheme described in [Gupta and Gao 2010]. It introduces a key generation resistant against coercion attacks. The idea is to incorporate skin conductance measurements into the cryptographic key generation. They experimentally show that the skin conductance measurement will help to reveal user's emotional states and recognize the attack as a stressful event (significantly different from the state when the keys were generated). This way, the generated keys include a dynamic component that can detect whether a user is forced to grant an access to the system.

10. CONCLUSION

We proposed a new biometric modality based on the human body's response to an electric square pulse signal. This biometric characteristic can serve an additional authentication mechanism in a PIN entry system, enhancing security of PIN entry with minimal extra user burden. The same biometric characteristic is applicable to continuous authentication. To this end, we proposed a continuous authentication mechanism on a secure terminal, which ensures user continuity, i.e., the user who started the session is the same one who is physically at the terminal keyboard throughout the session.

Through experiments with a proof-of-concept prototype we demonstrated that each human body exhibits a unique response to a signal pulse applied at the palm of one hand, and measured at the palm of the other. Using the prototype we could identify users in a matter of seconds. This identification mechanism integrates well with other established methods, e.g., PIN entry, to produce a reliable added security layer, either on a continuous basis or at login time.

We also focused our attention on how likely a legitimate user can be impersonated by an attacker using his own biometric data. We give average probabilities, as well as, experimental lower bounds found through simulations of worst-case scenarios.

REFERENCES

L.C.F. Araujo, Jr. Sucupira, L.H.R., M.G. Lizarraga, L.L. Ling, and J.B.T. Yabu-Uti. 2005. User authentication through typing biometrics features. *Signal Processing, IEEE Transactions on* 53, 2 (Feb 2005), 851–855. DOI:http://dx.doi.org/10.1109/TSP.2004.839903

Salil P Banerjee and Damon L Woodard. 2012. Biometric authentication and identification using keystroke dynamics: A survey. *Journal of Pattern Recognition Research* 7, 1 (2012), 116–139.

Claude Barral and Assia Tria. 2009. Fake Fingers in Fingerprint Recognition: Glycerin Supersedes Gelatin. In *Formal to Practical Security*, Vronique Cortier, Claude Kirchner, Mitsuhiro Okada, and Hideki Sakurada (Eds.). Lecture Notes in Computer Science, Vol. 5458. Springer, Berlin, 57–69. DOI:http://dx.doi.org/10.1007/978-3-642-02002-5_4

Saleh Bleha, Charles Slivinsky, and Bassam Hussien. 1990. Computer-access security systems using keystroke dynamics. *IEEE Transactions on pattern analysis and machine intelligence* 12, 12 (1990), 1217–1222.

Arman Boehm, Dongqu Chen, Mario Frank, Ling Huang, Cynthia Kuo, Tihomir Lolic, Ivan Martinovic, and Dawn Song. 2013. SAFE: Secure authentication with Face and Eyes. In *International Conference on Privacy and Security in Mobile Systems, (PRISMS)*. 1–8. DOI:http://dx.doi.org/10.1109/PRISMS.2013.6927175

Sungzoon Cho, Chigeun Han, Dae Hee Han, and Hyung il Kim. 2000. Web Based Keystroke Dynamics Identity Verification using Neural Network. *Journal of Organizational Computing and Electronic Commerce* 10 (2000), 295–307.

Nathan L. Clarke and Steven Furnell. 2007. Advanced user authentication for mobile devices. *Computers & Security* 26, 2 (2007), 109–119. DOI:http://dx.doi.org/10.1016/j.cose.2006.08.008

Cory Cornelius, Ronald Peterson, Joseph Skinner, Ryan Halter, and David Kotz. 2014. A Wearable System That Knows Who Wears It. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '14)*. 55–67. DOI:http://dx.doi.org/10.1145/2594368.2594369

Cory Cornelius, Jacob Sorber, Ronald A. Peterson, Joe Skinner, Ryan J. Halter, and David Kotz. 2012. Who Wears Me? Bioimpedance as a Passive Biometric. In *Proceedings of the 3rd USENIX Workshop on Health Security and Privacy (HealthSec 12)*, Carl A. Gunter and Zachary N. J. Peterson (Eds.).

Alexander De Luca, Alina Hang, Frederik Brudy, Christian Lindner, and Heinrich Hussmann. 2012. Touch Me Once and I Know It's You!: Implicit Authentication Based on Touch Screen Patterns. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 987–996. DOI:http://dx.doi.org/10.1145/2207676.2208544

I. Deutschmann, P. Nordstrom, and L. Nilsson. 2013. Continuous Authentication Using Behavioral Biometrics. *IT Professional* 15, 4 (July 2013), 12–15. DOI:http://dx.doi.org/10.1109/MITP.2013.50

Simon Eberz, Kasper Bonne Rasmussen, Vincent Lenders, and Ivan Martinovic. 2015. Preventing Lunchtime Attacks: Fighting Insider Threats With Eye Movement Biometrics. In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2014*.

Tom Fawcett. 2006. An introduction to ROC analysis. *Pattern Recognition Letters* 27, 8 (2006), 861–874. DOI:http://dx.doi.org/10.1016/j.patrec.2005.10.010

Tao Feng, Ziyi Liu, Kyeong-An Kwon, Weidong Shi, B. Carbunar, Yifei Jiang, and N. Nguyen. 2012. Continuous mobile authentication using touchscreen gestures. In *Homeland Security (HST), 2012 IEEE Conference on Technologies for*. 451–456. DOI:http://dx.doi.org/10.1109/THS.2012.6459891

Mario Frank, Ralf Biedert, Eugene Ma, Ivan Martinovic, and Dawn Song. 2013. Touchalytics: On the Applicability of Touchscreen Input as a Behavioral Biometric for Continuous Authentication. *IEEE Transactions on Information Forensics and Security* 8, 1 (2013), 136–148. DOI:http://dx.doi.org/10.1109/TIFS.2012.2225048

Javier Galbally, Arun Ross, Marta Gomez-Barrero, Julian Fierrez, and Javier Ortega-Garcia. 2012. From the Iriscode to the Iris: A New Vulnerability of Iris Recognition Systems. White paper. In *Briefings of the Black Hat Conference*.

Hugo Gamboa and Ana Fred. 2004. A behavioral biometric system based on human-computer interaction, In Biometric Technology for Human Identification. *Proceedings of SPIE* 5404 (2004), 381–392. DOI:http://dx.doi.org/10.1117/12.542625

R. Giot, M. El-Abed, and C. Rosenberger. 2009. Keystroke dynamics with low constraints SVM based passphrase enrollment. In *Biometrics: Theory, Applications, and Systems, 2009. BTAS '09. IEEE 3rd International Conference on*. 1–6. DOI:http://dx.doi.org/10.1109/BTAS.2009.5339028

Daniele Gunetti and Claudia Picardi. 2005. Keystroke analysis of free text. *ACM Transactions on Information and System Security (TISSEC)* 8, 3 (2005), 312–347.

Payas Gupta and Debin Gao. 2010. Fighting coercion attacks in key generation using skin conductance. In *Proceedings of the 19th USENIX Conference on Security (USENIX Security '10)*. USENIX Association, Berkeley, CA, USA, 30–30.

Sylvain Hocquet, J-Y Ramel, and Hubert Cardot. 2005. Fusion of methods for keystroke dynamic authentication. In *Fourth IEEE Workshop on Automatic Identification Advanced Technologies (AutoID'05)*. IEEE, 224–229.

Christian Holz and Marius Knaust. 2015. Biometric touch sensing: Seamlessly augmenting each touch with continuous authentication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. ACM, 303–312.

Information Technology Laboratory – National Institute of Standards and Technology. 2013. The Biometrics Resource Center. (2013).

Anil K. Jain, Arun Ross, and Karthik Nandakumar. 2011. *Introduction to Biometrics*. Springer.

Anil K. Jain, Arun Ross, and Sharath Pankanti. 2006. Biometrics: a tool for information security. *IEEE Transactions on Information Forensics and Security* 1, 2 (June 2006), 125–143. DOI:http://dx.doi.org/10.1109/TIFS.2006.873653

Zach Jorgensen and Ting Yu. 2011. On mouse dynamics as a behavioral biometric for authentication. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security (ASIACCS '11)*. 476–482. DOI:http://dx.doi.org/10.1145/1966913.1966983

Rick Joyce and Gopal Gupta. 1990. Identity authentication based on keystroke latencies. *Commun. ACM* 33, 2 (Feb. 1990), 168–176. DOI:http://dx.doi.org/10.1145/75577.75582

Kevin S. Killourhy. 2012. *A Scientific Understanding of Keystroke Dynamics*. Ph.D. Dissertation. Carnegie Mellon University, Pittsburgh, PA.

Lyra Nara. 2013. Hand Electrodes Brass (1 Pair). (2013). http://www.lyranara.com/hand-electrodes-brass-1-pair/

Orjan G Martinsen and Sverre Grimnes. 2011. *Bioimpedance and bioelectricity basics*. Academic press.

Fabian Monrose, Michael K. Reiter, and Susanne Wetzel. 1999. Password hardening based on keystroke dynamics. In *Proceedings of the 6th ACM conference on Computer and communications security (CCS '99)*. 73–82. DOI:http://dx.doi.org/10.1145/319709.319720

Youssef Nakkabi, Issa Traoré, and Ahmed Awad E. Ahmed. 2010. Improving mouse dynamics biometric performance using variance reduction via extractors with separate features. *IEEE Transactions on Systems, Man, and Cybernetics, Part A* 40, 6 (Nov. 2010), 1345–1353. DOI:http://dx.doi.org/10.1109/TSMCA.2010.2052602

National Science & Technology Council. 2006. Biometrics Frequently Asked Questions. (2006).

Minh Duc Nguyen and Quang Minh Bui. 2009. Your face is NOT your password: Face authentication bypassing - Lenovo - Asus - Toshiba. White paper. In *Briefings of the Black Hat Conference*.

Koichiro Niinuma and Anil K. Jain. 2010. Continuous user authentication using temporal information, In Technology for Human Identification VII. *Proceedings of SPIE* 7667 (2010), 76670L–76670L–11. DOI:http://dx.doi.org/10.1117/12.847886

Maja Pusara and Carla E. Brodley. 2004. User re-authentication via mouse movements. In *Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security (VizSEC/DMSEC '04)*. 1–8. DOI:http://dx.doi.org/10.1145/1029208.1029210

Kasper Bonne Rasmussen, Marc Roeschlin, Ivan Martinovic, and Gene Tsudik. 2014. Authentication Using Pulse-Response Biometrics. In *Proceedings of the 21st Annual Network and Distributed System Security Symposium*.

Kenneth Revett and Sérgio Tenreiro de Magalhães. 2010. Cognitive Biometrics: Challenges for the Future. In *Global Security, Safety, and Sustainability*, Sérgio Tenreiro de Magalhães, Hamid Jahankhani, and Ali G. Hessami (Eds.). Communications in Computer and Information Science, Vol. 92. Springer, 79–86. DOI:http://dx.doi.org/10.1007/978-3-642-15717-2_10

Sensible Vision Inc. 2013. Facial Recognition Provides Continuous System Security. (2013). http://www.sensiblevision.com/en-us/fastaccessanywhere/overview.aspx

R. Spillane. 1975. Keyboard Apparatus for Personal Identification. *IBM Technical Disclosure Bulletin* 17, 3346 (1975).

Pin Shen Teh, Andrew Beng Jin Teoh, and Shigang Yue. 2013. A survey of keystroke dynamics biometrics. *The Scientific World Journal* 2013 (2013).

Chee Meng Tey, Payas Gupta, and Debin Gao. 2013. I can be You: Questioning the use of Keystroke Dynamics as Biometrics. In *20th Annual Network and Distributed System Security Symposium, NDSS 2013, San Diego, California, USA, February 24-27, 2013*. http://internetsociety.org/doc/i-can-be-you-questioning-use-keystroke-dynamics-biometrics

Dat Tran, Wanli Ma, Girija Chetty, and Dharmendra Sharma. 2007. Fuzzy and Markov Models for Keystroke Biometrics Authentication. In *Proceedings of the 7th WSEAS International Conference on Simulation, Modelling and Optimization (SMO'07)*. World Scientific and Engineering Academy and Society (WSEAS), Stevens Point, Wisconsin, USA, 89–94. http://dl.acm.org/citation.cfm?id=1353862.1353878

David Umphress and Glen Williams. 1985. Identity verification through keyboard characteristics. *International Journal of Man-Machine Studies* 23, 3 (1985), 263 – 273. DOI:http://dx.doi.org/10.1016/S0020-7373(85)80036-5

VIRDI Biometric. 2009. How to make the fake fingerprints (by VIRDI). Video. (Feb. 2009). http://www.youtube.com/watch?v=-H71tyMupqk last accessed 03.08.2013.

John Woodward, Nicholas Orlans, and Peter Higgins. 2003. *Biometrics*. McGraw-Hill/Osborne.

Nan Zheng, Aaron Paloski, and Haining Wang. 2011. An efficient user verification system via mouse movements. In *Proceedings of the 18th ACM conference on Computer and communications security (CCS '11)*. 139–150. DOI:http://dx.doi.org/10.1145/2046707.2046725

Received January 2016; revised July 2016; accepted March 2017