

# Authentication Using Pulse-Response Biometrics

Ivan Martinovic, Kasper Rasmussen, Marc Roeschlin

University of Oxford

Department of Computer Science

{ivan.martinovic,kasper.rasmussen,marc.roeschlin}@cs.ox.ac.uk

Gene Tsudik

University of California, Irvine

Department of Computer Science

gene.tsudik@uci.edu

## ABSTRACT

We propose a new biometric based on the human body's response to an electric square pulse signal, called *pulse-response*. We explore how this biometric 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 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. Using a prototype setup, we show that users can be correctly identified, with high probability, in a matter of seconds. This identification mechanism integrates well with other established methods and offers a reliable 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 as a practical biometric. The results are very encouraging, achieving accuracies of 100% over a static data set, and 88% over a data set with samples taken over several weeks.

## 1. INTRODUCTION

Many modern access control systems augment the traditional two-factor authentication procedure (something you know and something you have) with a third factor: “something you are”, i.e., some form of 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 based on the human body's response to a square pulse signal. We consider two motivating scenarios:

The first is the traditional access control setting described above where the biometric is used as an additional layer of security when a user enters a PIN, e.g., into a bank ATM.

---

A full version of this paper was presented at the *Network and Distributed System Security (NDSS) Symposium 2014*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2008 ACM 0001-0782/08/0X00 ...\$5.00.

The pulse-response biometric 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 are often made of metal, which makes capturing pulse-response biometric 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*, e.g., verifying that the user, who securely logged in earlier, is the same person currently present at the keyboard. For this scenario, we need a mechanism that periodically samples one or more biometric. However, for obvious usability reasons, ideally this would be done *unobtrusively*. The pulse-response biometric is particularly well-suited for this setting. Assuming that it can be made from – or coated by – a conductive material, the keyboard would generate the pulse signal and measure response, while the user (remaining oblivious) is typing. The main idea is that the user's pulse-response is captured at login time and the identity of the person currently at the keyboard can be verified transparently, at the desired frequency.

To assess the efficacy and feasibility of the pulse-response biometric, we built a prototype platform that enables 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 ability and stability of this biometric over time. We also explored two systems that apply the pulse-response biometric to the two sample scenarios discussed above: one to unobtrusively capture the biometric as an additional layer of security when entering a PIN, and the other to implement continuous authentication.

## 2. BACKGROUND

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

### 2.1 Biometrics

The meaning of *biometric* varies depending on context. Throughout this paper we use it to denote a measurable biological (anatomical and physiological) or behavioral *characteristic* that can be used for automated recognition of individuals.

Usually, biometric measurements are divided into two categories, physiological and behavioural [7]. The former relies on the physiology of a person, such as fingerprints, facial features or DNA. Behavioral biometrics are based on user

behavior, such as keystroke timings, speech patterns, hand-writing characteristics, gait and many others.

Physiological biometrics can help identify an individual among large pool of candidates. However, there are some caveats. 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 [3]. Also, some facial recognition systems can be fooled by an appropriately-sized photo of a legitimate user.

Behavioral biometrics measure user actions over time, i.e., for each action, there must be a beginning, an end, and a duration. Consequently, behavioural biometrics indirectly measure characteristics of the human body. Behavioral biometrics are learned and, therefore, can be also re-learned. However, the consensus in the literature seems to be that after reaching a certain age, changes in behaviour become more difficult to achieve, even with specific and sustained effort [10]. Behavioural biometrics can therefore be regarded as valid means of identification, even though they are neither as unique nor as permanent as their physiological counterparts. In most cases, behavioural biometrics are used to discern a user from a small(er) pool of candidates. One 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 Biometric Authentication vs. Identification

Authentication refers to identify confirmation or verification. When a user claims a certain identity (e.g., by inserting a card into an ATM or entering a user ID into a terminal and then typing in a PIN or a password) authentication entails deciding whether the claim is correct. The goal of the biometric classifier is to compare the current sample to the known template for that user. The classifier returns the likelihood of a match. We refer to this as a  $1 : 1$  *comparison*.

Authentication differs from identification, where the current sample comes from an unknown user, and the job of the biometric classifier is to match it to a known sample. We refer to this a  $1 : n$  *comparison*. Identification is further divided into two types: open-set and closed-set. We say that an identification is closed-set, if it is known *a priori* that the user is in the classifier database, i.e., the classifier must choose the best match from a pool of candidates. Otherwise, identification is considered open-set.

## 2.3 Design Goals

When designing a new biometric system it is important to take into account lessons learned from past and current systems. Design goals for biometric systems can be found in the literature, e.g., [4]. Our goals include, but are not limited to:

*Universal*: Must be universally applicable, to the extent required by the application. It is important for the biometric to apply to everyone who is intended to use the system.

*Unique*: Must be unique within the target population. For example, measuring someone's height would not work as an identification mechanism on a large scale. At the same time, (adult) height alone *can* usually identify individual family members.

*Permanent*: Must remain consistent over the period of use. Very few biometrics will stay constant over a lifetime, e.g., face geometry, voice, gait and writing. However, as long as the biometric is consistent over the lifetime of the system, these biometrics work well.

*Unobtrusive*: If the user can be identified passively, without interference, the biometric is much more likely to be accepted.

*Difficult to circumvent*: Ideally, a user should be unable to change the biometric at all. At a minimum, a user must not be able to modify his biometric to match that of another user.

## 3. PULSE-RESPONSE BIOMETRIC

The pulse-response biometric works by applying a low voltage pulse signal to the palm of one hand and measuring the body's 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 is captured by measuring the response in the user's hand. This response is then 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 raw data that is then fed to a classifier. Working in the frequency domain eliminates any need for aligning the pulses when they are measured.

The main reason for the ability of this biometric to distinguish between users is due to subtle differences in body conductivity, 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, 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.

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

## 4. LIVENESS AND REPLAY

A common problem with many biometric systems is liveness detection, i.e., determining whether the biometric sample represents a "live" user or a replay. For example, 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 made of putty or even a severed finger. Similarly, a face recognition system would need to make sure that it is not being fooled by a user's photo or a 3-D replica.

In traditional biometric systems, liveness is 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. Also, a good 3-D model of a human head can still fool such measures.

Fingerprint scanners often include some protection against replay. This might be accomplished by detecting other characteristics normally associated with a live finger, e.g., temperature, or presence of sweat or skin oils. Such counter-measures make it more difficult to use skin-tight gloves or “cold dead fingers” to fool the biometric system. Still, replay remains a major challenge, especially for low-end fingerprint readers.

In the context of the pulse-response biometric, unlike fingerprints or face recognition, it is difficult (yet not impossible) to separate the biometric from the individual to whom it belongs. If the adversary manages to capture a user’s pulse-response on some compromised hardware, replaying it successfully would require specialized hardware that mimics the exact conductivity of the original user. We believe that this is feasible: the adversary can devise a contraption that consists of flat 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 significantly harder than in cases of facial recognition (where a photo suffices) or fingerprints, which are routinely left – and can be lifted from – numerous innocuous locations.

Finally, the real power of the pulse-response biometric is evident when used for continuous authentication (see Section 6). Here, the person physically uses a secure terminal and constantly touches the keyboard as part of routine work. Authentication happens on a continuous basis and it is not feasible 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, but the authentication system will see input from the keyboard without the expected pulse-response measurement to accompany it, and will lock the session.

## 5. COMBINING PIN ENTRY WITH BIOMETRIC CAPTURE

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

### 5.1 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 the pulse-response classifier with probability higher than that found in our experiments described later in this paper.

We also assume that the ATM is equipped with a modified authentication module which, besides verifying the PIN, captures the pulse-response biometric and determines the likelihood of the measured response corresponding to the user identified by the inserted ATM card and the just-entered PIN. We assume that the ATM has access to a database of valid users, either locally or over a network. Alternatively, the user’s ATM card can contain data 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.

### 5.2 PIN Entry Scheme

The ATM has to determine whether data sampled from the user while entering the PIN is consistent with that stored 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 samples in the database to produce a match. Using our prototype, we can make such decisions with high confidence; see Section 7.4.

Before discussing security of the pulse-response PIN entry system, we check whether it meets the design goals.

*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, 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 check. This would allow our approach to gracefully degrade to a generic PIN entry system.

*Unique and Permanent.* In Section 7.4 we show that our prototype can determine, with high probability, whether a subject matches a specific pulse-response. Thus, it is extremely 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.* The proposed 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 hand on a metal pad. There can even be two such pads accommodating both left- and right-handed people. Also, 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 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 very hard to test precisely, assuming that the adversary is unaware of the target user’s pulse-response measurements, the task seems very difficult, if not impossible.

### 5.3 Security of PIN Entry Scheme

The additional layer of security provided by the pulse-response biometric 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:

$$P_{break} = P_{guess} \cdot P_{forge}$$

where  $P_{guess}$  is the probability of the adversary correctly guessing the PIN and  $P_{forge}$  is the average probability that the adversary can fool the classifier. We model this as the false positive rate divided by the number of users. The false positive rate, i.e., when an adversary is incorrectly classified as an authorized user, is the complement of specificity [9]. In Section 7.4, we determine specificity to be 88% and thus  $P_{forge} = (1 - 0.88)$  on average.

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

$$P_{break} = \frac{(1 - 0.88)t}{10^n}$$

For example, if the adversary is allowed 3 guesses with a 4-digit pin,  $P_{break} = 3.6 \cdot 10^{-5}$ , 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” and do not involve the adversary guessing the PIN. If we assume that the adversary already knows the PIN,  $P_{break} = 12\%$  with our system, as opposed to 100% without it.

## 6. CONTINUOUS AUTHENTICATION

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

### 6.1 System and Adversary Models

We continue using the example for continuous authentication introduced in Section 1. It entails a secure terminal where authorized users can login and access sensitive data.

The system consists of a terminal with a special keyboard that sends out pulse signals and captures the pulse-response biometric. 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.

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. We assume that the adversary at the keyboard has full access to the active session. The goal of our system 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

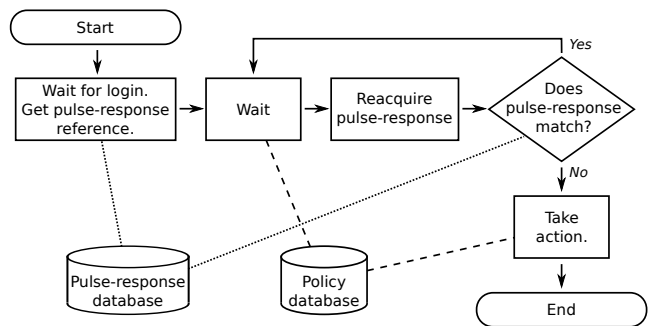


Figure 1: Flowchart of the *Continuous Authentication Process* decision procedure.

consists of a software process that manages initial login and frequency of periodic reacquisition of the biometric. 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.

### 6.2 Continuous Authentication Scheme

At login time, CAP measures and records the initial pulse-response biometric of the authorized user. Periodically, e.g., every few seconds, CAP reacquires the biometric 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 1 shows a sample CAP decision flowchart.

The envisaged continuous authentication system can be useful for training (e.g., corporate) users to adopt security-conscious behaviour. For example, users can be motivated to behave securely whenever they leave a secure terminal, e.g., by getting a warning every time they forget to log out, and/or allow someone else to take over a secure session.

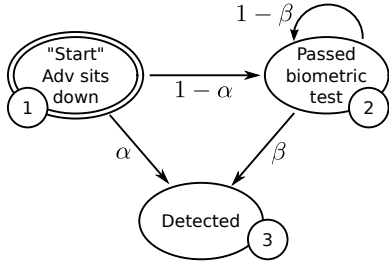
Before considering the security of the continuous authentication system, we look back at the design goals.

*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 7.4, we show that our prototype can match a pulse-response to previous samples (taken immediately beforehand) with 100% accuracy. 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.* Users do not need to modify their behaviour at all when using the continuous authentication system. Thus, user burden is minimal.

*Difficult to Circumvent.* With a true positive rate of 100% it is unlikely that the adversary can manage to continuously fool the classifier. Even if the adversary happens to have a pulse-response biometric similar to the original user, it must evade the classifier on a continuous basis. We explore this further in the security analysis section below.



**Figure 2: Markov model of the continuous authentication detection probability. States are numbered 1 to 3 for easy reference in text.**

### 6.3 Security

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. This results in a worst-case scenario and the detection probability is a lower bound on security provided by the continuous authentication system.

An important measure of security is the detection time—the number of times biometric acquisition is performed between the adversary’s initial appearance and detection. Obviously, longer inter-acquisition intervals imply slower collection of measurements and subsequent detection of adversarial presence.

We model the probability of detecting an adversary using two static probabilities derived from our experiments—an initial probability  $\alpha$  and a steady state probability  $\beta$ . A more detailed model with several intermediate decreasing probabilities could be constructed but this simple model fits quite well with our experiments.

The probability  $\alpha$  is the probability that the adversary is detected immediately, i.e., the very first time when his pulse-response is measured. However, if the adversary’s biometric is very close to that of the original user, the adversary might not be detected every time biometric capture is performed. This is because the biometric is subject to measurement noise and the measurements from an individual form a distribution around the “fingerprint” of that user. If the adversary manages to fool the classifier once, it must be because its biometric is close to that of the original user. Thus, the adversary’s subsequent detection probability must be lower:

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

We call this decreased probability  $\beta$ . The probabilities  $\alpha$  and  $\beta$  are approximations that model how similar two individuals are, i.e., how well their probability distributions overlap in about 100 dimensions. Using  $\alpha$  and  $\beta$  we build a Markov model, shown in Figure 2, 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  $\alpha$  or *not* detected, with probability  $1 - \alpha$ . In the latter case, its pulse-response biometric must be close the original user’s. Thus,  $\beta$  is used for the subsequent rounds. In each later round, the adversary is either detected with probability  $\beta$  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  $\rho$  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  $\rho P$ . Probabilities for each subsequent round are determined via another multiplication by  $P$ . The probabilities of being in each state after  $i$  rounds (state transitions), is therefore:

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

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 2). The probability of being in state 2, i.e., to escape detection for  $i$  rounds, is given by the second element of  $\rho$ :  $(1 - \alpha)(1 - \beta)^{i-1}$ . The probability of detection is thus:  $1 - (1 - \alpha)(1 - \beta)^{i-1}$ .

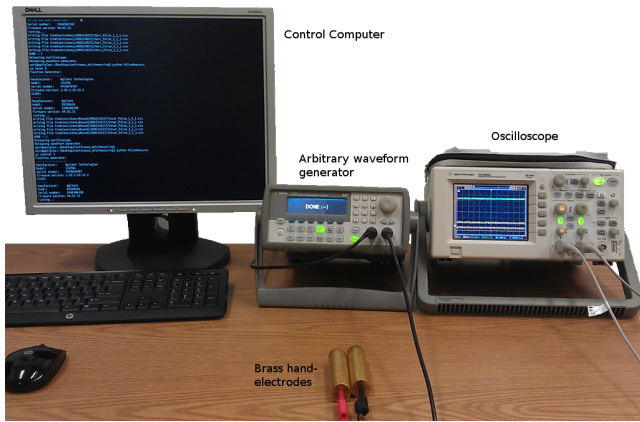
$\alpha$  roughly corresponds to the *sensitivity* of the classifier, i.e., the true positive rate reported in Section 7. We use 99% (rather than the 100% found in our experiments) in order to model the possibility of making a classification mistake at this point. We do not have enough data to say with absolute certainty if this is valid for very large populations, but we continue under the assumption that our data is representative.  $\beta$  is harder to estimate but we set  $\beta = .3$  based on numbers from our experiments in Section 7.4. Using these values there is a 99.96% chance of detecting the adversary after 10 rounds. This grows to 99.9999997% after 50 rounds. Thus, not surprisingly, acquisition frequency determines the time to detect the adversary.

What the very high 99.999+% detection probability is really saying is that, if you just test enough times, the authentication will eventually fail. It matches very well with our experiments and it is true even for a legitimate user (although much less frequently). For this reason we need a way to handle false negatives.

### 6.4 Handling False Negatives

False negatives refer to incorrect detection of adversarial presence. If the biometric 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 negatives 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. For example, allowing one event every 100 rounds corresponds to a false negative rate of 1%. Another option is to integrate a less user-friendly (less transparent) biometric 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.



**Figure 3: Proof-of-concept measurement setup.** The test subject holds two brass electrode handles 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.

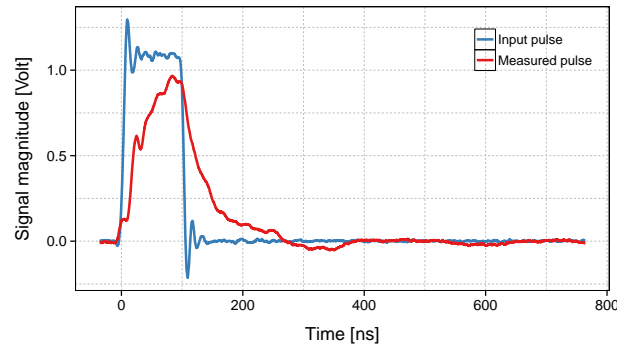
Yet another alternative is the gradual ramp up of the severity of actions taken by the continuous authentication process, for each successive detection event. For the first time, displaying a warning might be the most appropriate action. If detection re-occurs, more and more severe actions can be taken. It is very unlikely, with a reasonably low false negative rate, to have multiple consecutive adversary detection events if the original user is still at the terminal. Although the false positive rates we achieve are quite low, they could certainly be improved with a more advanced biometrics capture system. In conjunction with a sensible policy, our continuous authentication system might be appropriate for any organization with high security requirements.

## 7. EXPERIMENTS

Starting out with the hypothesis that the biometric measurement varies depending on the frequency of the signal transmitted through the human body, we rigorously experimented with various frequencies, voltage levels and waveforms. We also assessed several classification algorithms. Our experiments suggested the choice of 100ns long square pulses at 1 volt as the input signal (see Figure 4) and *Support Vector Machines (SVM)* for classifying samples. Hence, the name *pulse-response* biometric. Complete analysis can be found in the full version of this paper [?].

### 7.1 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 3 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.



**Figure 4: Input and output waveforms.** One measurement consists of 4,000 samples with the rate of 500 MSa/s.

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.

### 7.2 Ethics and User Safety

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.

### 7.3 Biometric Capture Procedure

Each subject followed a specific procedure during the biometric measurement process to ensure that only minimal noise is introduced into the measured data. 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 consisted of 22 samples for each subject, taken in one measuring session, i.e., at one point in time. The second included 25 samples per test person, obtained in five different sessions, over time. This was done to assess stability of the biometric over time.

The subject population included both males and females between the ages of 24 and 38. We sampled all test subjects at different times during the day over the course of several weeks. We tried to sample subjects in order to end up with sampling conditions as diverse as possible, for each subject. The interval between measurement sessions for the same subject was varied between several hours and several weeks. This was done in order to try to eliminate any effects of sampling at a specific time of the day.

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 4 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



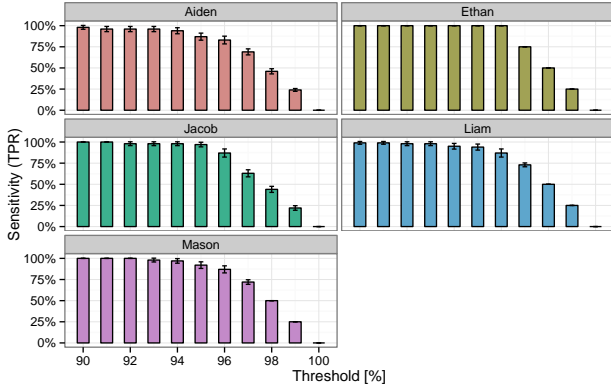


Figure 5: True positive rate for each test subject for the authentication classifier fed with the data sampled over time. Error bars show 95% confidence interval. The  $x$ -axis reflects the discrimination threshold for assigning the classifier’s prediction output to a *positive* or a *negative*.

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.

## 7.4 Results

We present two different classifiers: one for authentication and one for identification. The former is based on support vector machines (SVM) and verifies a 1 : 1 match between a sample from an unknown person and that of a requested person. The identification classifier, also based on SVM, verifies a 1 :  $n$  match between a sample of a known person against all samples in a database. The identification classifier is of a closed-set variety. Section 2 provides a more detailed description of open- and closed-set classifiers.

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

### 7.4.1 Authentication Classifier

Figure 5 shows the distinguishing potential of the authentication classifier applied to a data set collected over several weeks. Each bar shows the classifier’s performance for different threshold levels, for each of the test subjects. The threshold is a measure of assurance of correct identification. If a low false positive rate is acceptable, better sensitivity can be achieved. The classifier’s performance is measured using 5-fold cross-validation to ensure statistical robustness. The figure shows that all subjects are recognized with a very high probability, as the true positive rate confirms.

Applying the authentication classifier to the single-session data set yields even better performance figures (see the full version of this paper [?]). For example, 10% false positives allow us to achieve sensitivity of almost 100%.

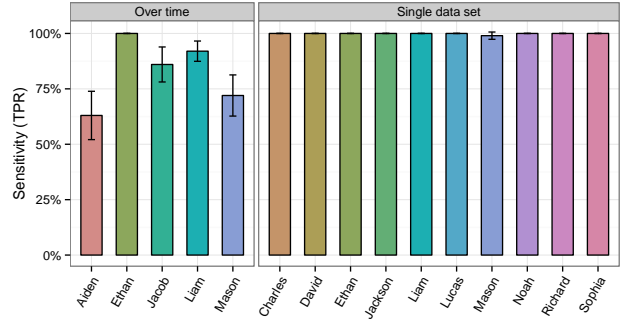


Figure 6: Identification classifier results. The true positive rate for each test subject is obtained by applying 5 times stratified 5-fold cross-validation. Error bars show 95% confidence interval.

	TP	FP	TN	FN	in [%]		
					Sensitivity	Specificity	Accuracy
Authentication							
– Single set	2.0	0.0	18.0	0.0	100	100	100
– Over time	4.4	2.4	17.6	0.6	88	88	88
Identification							
– Single set	2.0	0.0	18.0	0.0	100	100	100
– Over time	3.4	1.6	18.4	1.6	68	92	87.2

Table 1: Summary of results for authentication and identification classifiers, averaged over all users. All performance figures have been assessed on the basis of test data not involved in any development or training phase of the classifiers. Values for true/false positives/negatives are at the equal error rate of  $EER = 0.00$  for the single data set and  $EER = 1.12$  over time.

### 7.4.2 Identification Classifier

Identification is a multi-class classification problem. Our classifier consists of multiple SVMs and follows a one-against-one approach (aggregation by voting). Due to this increased complexity a slight drop in performance is expected, in comparison to authentication, which is a binary classification task.

Results obtained from the identification classifier over the two data sets are shown in Figure 6. Even with increased complexity, the identification classifier performs very well on both data sets. The single-session data set contains ten people and the goal of the classifier is to identify each person as accurately as possible. There is a slight decrease in performance for the data set containing samples taken several weeks apart. The reason for this decrease is that samples 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, and a number of other factors.

Table 1 summarizes results for the two classifiers. Both classifiers can be tuned by selecting a specific false positive rate. For example, in a continuous authentication application, where false negatives are of greater concern, classifiers can

be tuned to a lower false negative rate, by accepting a higher false positive rate.

## 8. RELATED WORK

The full version of this paper has a detailed survey of related work [?]. In this version we provide a brief overview.

Biometrics, as a means of recognizing an individual using physiological or behavioural traits, has been an active research area for many years. A comprehensive survey of conventional physiological biometrics can be found in [5]. While physiological biometrics tend to be relatively stable over time, they are sensitive to deception attacks, e.g., mock fingers [1]. In contrast, behavioural biometrics are much harder to circumvent. However, the performance of behavioral biometric systems is usually worse and can require re-calibration due to normal variations in human behaviour. Initial results on behavioral biometrics were focused on typing and mouse movements, e.g., [8]. Keystroke dynamics became quite popular [6], as a means to augment password authentication in manner similar to our PIN-entry scenario.

The result most closely related to our work is [2], where bioimpedance is used as a biometric: a wearable wrist sensor passively recognizes its wearers based on the body's unique response to the alternating current of different frequencies. Experiments in [2] were conducted in a family-sized setting and show a recognition rate of 90% when measurements are augmented with hand geometry. The pulse-response biometric proposed in this paper solves a different problem but it also uses the body's response to a signal. It achieves a recognition rate of 100% when samples are taken in one session and 88% when samples are taken weeks apart (no augmentation is required in both cases).

## 9. CONCLUSION

We proposed a new biometric based on the human body's response to an electric square pulse signal. This biometric can serve an additional authentication mechanism in a PIN entry system, enhancing security of PIN entry with minimal extra user burden. The same biometric is applicable to continuous authentication. To this end, we designed 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 applied at the palm of one hand, and measured at the palm of the other. Using the prototype we could identify users – with high probability – 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.

## 10. REFERENCES

- [1] C. Barral and A. Tria. Fake fingers in fingerprint recognition: Glycerin supersedes gelatin. In V. Cortier, C. Kirchner, M. Okada, and H. Sakurada, editors, *Formal to Practical Security*, volume 5458 of *Lecture Notes in Computer Science*, pages 57–69. Springer Berlin Heidelberg, 2009.
- [2] C. Cornelius, J. Sorber, R. Peterson, J. Skinner, R. Halter, and D. Kotz. Who wears me? bioimpedance as a passive biometric. In *Proceedings of the USENIX Workshop on Health Security and Privacy*, August 2012.
- [3] N. S. . T. Council. Biometrics frequently asked questions, 2006.
- [4] A. Jain, A. Ross, and K. Nandakumar. *Introduction to Biometrics*. SpringerLink : Bücher. Springer, 2011.
- [5] A. Jain, A. Ross, and S. Pankanti. Biometrics: a tool for information security. *Information Forensics and Security, IEEE Transactions on*, 1(2):125 – 143, June 2006.
- [6] F. Monrose, M. K. Reiter, and S. Wetzal. Password hardening based on keystroke dynamics. In *Proceedings of the 6th ACM conference on Computer and communications security, CCS '99*, pages 73–82, New York, NY, USA, 1999. ACM.
- [7] I. T. L. N. I. of Standards and Technology. The biometrics resource center, 2013.
- [8] R. Spillane. Keyboard apparatus for personal identification. *IBM Technical Disclosure Bulletin*, 17(3346), 1975.
- [9] Wikipedia. Sensitivity and specificity, 2013.
- [10] J. Woodward, N. Orleans, and P. Higgins. *Biometrics*. RSA Press Series. McGraw-Hill/Osborne, 2003.