Broken Hearted: How To Attack ECG Biometrics

Simon Eberz
University of Oxford
simon.eberz@cs.ox.ac.uk

Nicola Paoletti
University of Oxford
nicola.paoletti@cs.ox.ac.uk

Marc Roeschlin
University of Oxford
marc.roeschlin@cs.ox.ac.uk

Andrea Patané
University of Catania
andrea.patane@eng.ox.ac.uk

Marta Kwiatkowska
University of Oxford
marta.kwiatkowska@cs.ox.ac.uk

Ivan Martinovic
University of Oxford
ivan.martinovic@cs.ox.ac.uk

Abstract—In this work we present a systematic presentation attack against ECG biometrics. We demonstrate the attack's effectiveness using the Nymi Band, a wrist band that uses electrocardiography (ECG) as a biometric to authenticate the wearer. We instantiate the attack using a hardware-based Arbitrary Waveform Generator (AWG), an AWG software using a computer sound card, and the playback of ECG signals encoded as .wav files using an off-the-shelf audio player. In two sets of experiments we collect data from a total of 41 participants using a variety of ECG monitors, including a medical monitor, a smartphone-based mobile monitor and the Nymi Band itself.

We use the first dataset to understand the statistical differences in biometric features that arise from using different measurement devices and modes. Such differences are addressed through the automated derivation of so-called mapping functions, whose purpose is to transform ECG signals from any device in order to resemble the morphology of the signals recorded with the Nymi Band.

As part of our second dataset, we enroll users into the Nymi Band and test whether data from any of our sources can be used for a signal injection attack. Using data collected directly on the Nymi Band we achieve a success rate of 81%. When only using data gathered on other devices, this rate decreases to 43% when using raw data, and 62% after applying the mapping function. While we demonstrate the attack on the Nymi Band, we expect other ECG-based authentication systems to most likely suffer from the same, fundamental weaknesses.

I. INTRODUCTION

Passwords are the most prevalent mode of authentication in many environments, including local workstations and on the web. Despite their widespread use, they suffer from a number of weaknesses [12], [13], [3]. Most notably, these include users choosing weak passwords [15] and frequent password-reuse [23]. Besides these weaknesses, the increasing popularity of mobile and wearable devices gives rise to another challenge: the lack of input devices to enter passwords (e.g., smart watches might only provide few buttons for user input). Biometric recognition has become a popular approach to tackle these limitations. Unlike passwords, which rely on the user knowing something, biometrics make use of either distinctive physiological properties or behavior. The former includes fingerprints, iris patterns and DNA, while keystroke dynamics, touchscreen input, mouse movements and eye movements have been proposed for the latter. Electrocardiography (ECG) records the electrical activity of the heart over time. While ECG is typically recorded in a hospital using 10 electrodes placed on the patient's skin, it can also be measured using two electrodes, thus making it feasible to record with wearable devices. Due to its distinctiveness and universality (every living human has a heartbeat that can be measured), ECG as a biometric has attracted considerable attention in recent years. Unlike most behavioral biometric systems, which exist mostly as research projects, ECG-based biometrics have resulted in a successful commercial product, the Nymi Band[1].

In this paper, we provide a systematic attack against ECG biometrics and demonstrate its effectiveness by applying it to the Nymi Band. In order to this end, we first demonstrate our capability of spoofing arbitrary ECG signals. To present the spoofed signals we use three different devices, two Arbitrary Waveform Generators (AWG), one software- and one hardware-based, as well as the audio playback of ECG signals encoded as .wav files using an off-the-shelf audio player. As such, the technological barriers for the attacker are extremely low. We collect ECG data from a total of 41 users and 5 devices. The data shows that the morphology of the ECG signal depends greatly on the device that was used. This difference constitutes a major challenge, as a signal collected on one device can not easily be used to carry out an attack on another. We tackle this challenge by using the first half of our dataset to devise and train a so-called mapping function. The purpose of this function is to transform an ECG signal collected on one device such that its morphology matches that of another. The method to generate these functions is general and can even be used for other biometrics. After enrolling users in the Nymi Band we use the mapping function to impersonate the user by presenting the (transformed) ECG signals collected on different devices.

The rest of the paper is organized as follows: Section II gives background information about ECG and ECG biometrics. Section III describes our hardware setup and the approach used to spoof arbitrary ECGs. We outline our experimental design and data collection in Section IV. In Section V we describe the generation and training of the mapping function and outline

---

[1]http://www.nymi.com

the results of carrying out the attack in Section VI. Section VII presents related work, we discuss possible countermeasures to our attack in Section VIII and conclude the paper in Section IX.

II. Background

A. Electrocardiography

The electrocardiogram (ECG) is a measurement of the electrical activity of the heart. It is acquired through electrodes placed on the patient's skin, which are used to capture voltage changes due to depolarization and repolarization of cardiac cells, respectively provoking contraction and relaxation of the cardiac muscle. The ECG is commonly used in clinical practice for its crucial diagnostic capabilities [14]. In addition, the present availability of low-cost ECG sensors has opened to numerous applications in the area of wearable devices and fitness monitoring [18], leading to pervasive acquisition of ECG data.

Figure 1 shows an example ECG for one cardiac cycle, together with the duration and amplitude features typically extracted for authentication purposes (detailed in Section II-B). It comprises five main waves, P, Q, R, S and T, which map to specific heart events: the *P wave* indicates activation of the atria (the upper heart chambers); the *QRS complex* corresponds to the activation of the ventricles (the lower chambers); and the T wave indicates ventricular repolarization.

Most of ECG recording systems are based on the so-called Einthoven's lead system, where each lead records the difference of potential between two electrodes. Einthoven's leads consist of:

$$\text{Lead I:} \quad \text{Lead II:} \quad \text{Lead III:}$$
$$V_I = \Phi_{LA} - \Phi_{RA} \quad V_{II} = \Phi_{LF} - \Phi_{RA} \quad V_{III} = \Phi_{LF} - \Phi_{LA}$$

where V_i is the voltage of lead i and Φ_j, with $j \in \{LA, RA, LF\}$, is the potential at the left arm (LA), right arm (RA) and left foot (LF), respectively. In particular, the standard 12-lead ECG used in clinical settings is an extension of the Einthoven's 3-lead system based on using seven additional electrodes placed on the chest. Nevertheless, simpler 1-lead ECG recording systems are increasingly being used in the context of personal ECG monitoring and wearables.

B. ECG Biometrics

Driven by the distinctiveness and universality of ECG, the body of work in this field has been steadily growing over the past few years. Recent surveys of systems based on ECG-based biometrics can be found in [24], [8], [1]. The most striking difference between approaches lies in the biometric features used. The first class of methods is based on *time domain feature extraction*, and work by detecting the so-called fiducial points, i.e., location, amplitude and width of the main ECG waves, as shown in Figure 1. Some biometric systems also consider the ST segment, that is, the length of the isoeletric segment between the S and T waves, as well as the slope of waves. In Figure 1, the wave slope is accounted for through the extraction of left and right components of its width. In addition to the above intra-beat morphological features, inter-beat features such as Heart Rate Variability and beat patterns (represented by the RR intervals) can well reflect the specific characteristics of the subject.


Fig. 1: Example electrocardiogram and corresponding time-domain features for ECG-based biometrics. Top: duration features given by inter-peak distances. Right: amplitude features. For each wave, we also consider its width at half amplitude (grey solid lines). To account for asymmetric curves, the width of each wave is split into left and right components (L and R segments shown in P and T waves), i.e., before and after the wave peak.

The second class of methods use *frequency domain feature extraction*, meaning that features are obtained after converting the ECG signal in the frequency domain. Examples include application of wavelet decomposition and Fourier Transform. In alternative to time and frequency domain methods, some biometric systems employ statistical approaches for computing the distance between enrolment ECG and recognition ECG directly at the signal level, or analysis of the ECG's trajectory in the phase space.

Despite the numerous advantages of ECG-based biometrics, they also suffer from some limitations, mostly addressed through the use of multimodal biometrics. These are related to the time-variant nature of the ECG, which is affected by physical activity, emotional stress and minor random disturbances like ectopic beats.

C. The Nymi Band

In this subsection we will outline the capabilities and system design of the Nymi Band.

The Nymi Band (see Figure 2) is a wristband that incorporates an ECG sensor with two electrodes. The bottom electrode constantly touches the user's wrist while the band is worn. In order to allow ECG measurements (most commonly for enrollment and authentication) the user touches the top electrode with the index finger of their other hand. As such, the signal morphology can be expected to be similar to Lead I of a medical ECG (which measures the potential difference between the left and right arm). Besides the actual band, the Nymi ecosystem consists of the Nymi Companion App (NCA) and Nymi Enabled Applications (NEAs). The NCA is provided

Fig. 2: The Nymi Band

as an app that runs on the user's smartphone or tablet. The NCA performs two main functions, enrollment and activation.

During *enrollment*, the Nymi Band is paired to an NCA. The correct pairing is confirmed by displaying a pattern on the Nymi Band which the user has to verify against a pattern shown by the NCA (similar to the numerical codes used in Bluetooth device pairing). The Nymi Band and NCA then agree on a shared key that binds the Nymi Band to this NCA. Following pairing, the user is prompted to touch the band's top electrode with his index finger, after which their ECG is measured until a specific amount of ECG data of sufficient quality is captured. The resulting biometric template is then encrypted and stored by the NCA on the phone or tablet. Besides the shared secret, no information is stored on the band at this time.

Activation is performed when the Nymi Band is taken off and put back on again. Specifically, this event is detected by the contact between two pins on the inside of the buckle being interrupted (see Figure 2). As such, the Nymi Band does not truly perform continuous authentication in the biometric sense, but authenticates the user once and then detects a possible change in user identity through the band being taken off. The activation process is started by the user selecting the appropriate action in the NCA, after which they can choose to perform ECG authentication, or to use their backup password. If they choose ECG, they are again prompted to touch the top electrode to begin ECG measurement. Unlike enrollment, which runs until a certain number of seconds of valid ECG data is collected, activation runs until the NCA is sufficiently convinced of the wearer's identity. Once one or several heartbeats are observed that match the owner's template, the Nymi Band is put into activated mode by the NCA. If no matching heartbeats are observed after 60 seconds, the user is automatically rejected.

Once the Nymi Band is activated, it can be paired with NEAs. Examples of NEAs include desktop computers (that can then be unlocked without using a password), wearable devices like smart watches and even more complex systems like cars. At the time of writing, the Nymi Band is being trialled for contactless payments. Initially, the band is paired with the NEA through a process similar to regular Bluetooth pairing. The Nymi Band displays a pattern using the five LEDs (leading to only 32 possible combinations), which the user is meant to confirm before proceeding. The Nymi Band and the NEA then use a Diffie-Hellman key exchange to negotiate

a shared key, which is stored directly on the band. After pairing, the possession of the shared key (i.e., the presence of the unlocked band) can then be confirmed using a standard challenge-response protocol.

There is one additional capability of NEAs that is relevant to the remainder of the paper: The Nymi SDK grants NEAs direct access to the band's ECG sensor. Once the band and an NEA are paired, the NEA can request the collection of an arbitrary amount of raw ECG data. While this data collection does not have to be explicitly approved by the user, the sensor design requires the user to touch the top electrode with their finger, thus making covert data collection virtually impossible. It is noteworthy that this functionality has been removed from the official SDK from version 2.0 onwards.

The Nymi Band's threat model is described in the Nymi Whitepaper. The band is designed as a three-factor authentication system. In order to communicate with NEAs, an attacker has to be in possession of the Nymi Band and the NCA (typically the user's phone) and be able to bypass the biometric authentication. It is noteworthy that the latter, while not explicitly stated in the Whitepaper, can also be achieved by using the user's backup password (e.g., through guessing a weak password or social engineering). This is particularly dangerous, as the presence of a second authentication factor often leads to users choosing weaker passwords [26]. In terms of bypassing ECG authentication (rather than using a password), the Nymi Whitepaper claims that

> *"There is currently no known means of falsifying an ECG waveform and presenting it to a biometric recognition system."*

In the following sections we will investigate the validity of this claim.

III. SPOOFING ECG SIGNALS

In this section, we show that fake ECG signals can be injected into ECG enabled recognition systems. We start out with the hypothesis that captured ECG measurements can be reproduced at the biometric sensors without the benign user having to be present.

A. Motivation

Like any other physiological trait, ECG signals can be captured and (digitally) stored for an indefinite amount of time as the signals are relatively immutable. Biometric samples from physiological traits do not lose validity and, if the fidelity of the stored signal is sufficiently high, it is possible to physically reproduce the actual biometric signal at a later time. This process does not require the individual from whom the biometric measurements originate to be present.

In ECG recognition, biometric readings are usually acquired with the help of an electrocardiograph, which works by measuring the minute voltage differences of the human heart over time. With today's technologies in signal synthesis and digital to analog conversion, artificially creating electrical signals that exactly represent stored ECG signals is feasible.

While forging ECG signals is not a concern in the medical domain, it is potentially problematic for ECG-based authentication systems. If a biometric system does not feature an agent or

overseer — or other provisions against someone not using the biometric sensors as intended — it is susceptible to so-called presentation attacks. In a presentation attack the attacker tries to spoof the biometric sensors with an artefact or contraption. In case of ECG based recognition, the attacker would have to fake the (time-dependent) voltage levels at the electrodes interfacing the user with the help of an electrical device that outputs an ECG signal.

In the remainder of this section, we show that it is indeed possible to replay previously captured ECG signals. To that end, we built three hardware contraptions of varying degrees of sophistication that successfully create and inject ECG signals into the sensing electrodes of a biometric system based on ECG. We test our contraptions using the example of the Nymi Band. In order to estimate difficulty and likelihood of a presentation attack on ECG recognition, we additionally evaluate our injection methods along the following non-technical dimensions:

- **Cost**: What is the overall cost for building the contraption and executing the injection? Although high cost does not deter every attacker, it can discourage less determined ones.

- **Knowledge**: Is expert knowledge required to build and use the contraption or can it be put together following simple instructions?

- **Size**: Physical size is a very important factor. If the contraption used for signal injection is sufficiently small, an attacker can covertly spoof the biometric sensors and might even circumvent a guarded biometric system.

- **Signal quality**: We quantize and compare the resulting signal quality of each approach when applied to injecting ECG signals into the Nymi Band. Obviously, signal quality directly correlates with the probability of success for a presentation attack. The conversion from the (stored) biometric data to the physical biometric signal should introduce as little noise as possible.

B. Hardware Considerations

We build three contraptions that allow us to forge fake ECG signals. Since all approaches inject the resulting signal at the sensing electrodes of the Nymi Band, we design a contrivance that allows us to interface those electrodes in a convenient way. As described in the previous section, the Nymi Band has a locking mechanism, which deauthenticates the wearer of the band immediately if the band is taken off. When the band is closed, one of the two sensing electrodes is located on the inside of the band and faces the wearer's wrist whereas the other electrode is accommodated at the outside of the band, available to be touched with a finger of the other hand. This way, the electrical circuit is closed and the ECG measurement can start. If the wearer of the band is recognized, the band goes to and remains in "authenticated" state as long as the band is worn.

For ease of use and to allow many successive injection attempts without opening and closing the Nymi Band, we modified a genuine charging cable that is included in the delivery of the band. As can be seen in Figure 3, the part of the


Fig. 3: Modified charger lead. The two pins on the left are shorted out to put the Nymi Band into "closed" state as soon as the modified charging lead is attached to the band.

charging lead that interfaces with the band has two shorted-out pins to let the band think it is closed and conveniently exposes one of the sensing electrodes in a separate wire.

1) Hardware Waveform Generator: Our first and the most obvious approach to artificially create an ECG trace is to use an arbitrary waveform generator. The purpose of waveform generators is to generate electrical waveforms for testing and analyzing electronic devices. The generated signal is injected into the device under test while the device's (electrical) behavior is observed and analyzed. The waveform itself is defined as a time series of voltage levels, based on which the waveform generator produces the corresponding signal. The resulting signal exactly matches the predefined voltage levels at the given points on the time axis and interpolates values in between.

We use a Rigol DG4062 Arbitrary Waveform Generator that is capable of generating signals of up to 60MHz at a sampling rate of 500 Mega-Samples per second. These specifications enable us to transform stored ECG signals to their physical counterpart with high accuracy. Electrocardiographs commonly operate at a sampling rate of less than 500 samples per second when acquiring an ECG trace. This means that higher frequency components, i.e., more than 250Hz, can not be registered and hence are not part of the measured signal. Such a frequency limited signal can easily be synthesized by most of today's off-the-shelf hardware waveform generators.

Our signal generator's two output leads are directly connected to the electrodes of the Nymi Band. In order to optimally match electrical impedances between signal generator and the electrodes of the band, we inject the signal through a 75Ω coaxial cable (see Figure 4). Electrocardiographs most often feature an instrumentation amplifier with high input impedance as the first step in the signal acquisition pipeline. The Nymi Band does not differ in that regard and requires a relatively low impedance input.

We wrote a software library that loads stored ECG signals directly into the memory of the Rigol DG4062 Arbitrary Waveform Generator, sets the necessary parameters and starts/stops the signal generation. The program code is available upon request.

Fig. 4: Arbitrary waveform generator connected to the Nymi Band via the modified charging lead. The negative output of the waveform generator is clamped to the electrode facing the wrist and the positive output is attached to the second electrode of the band using the modified charging lead.


Fig. 5: A laptop is connected to the Nymi Band via the modified charging lead. Setup is analogous to the configuration involving the hardware waveform generator, apart from the coaxial cable being plugged into the audio output port of a laptop. The laptop either runs a software waveform generator or is used to play back an ECG signal that is encoded in an audio file. The laptop might be replaced through any electronic device with audio playback capability.

2) Software Waveform Generator: Nowadays, almost every personal electronic device, be it mobile or stationary, possesses a dedicated sound card or integrated sound functionality to facilitate analog audio output. Audio signals are an electrical representation of sound, i.e., a mechanical wave that propagates through a medium. Thus, sound cards need to be able to output relatively high-frequency signals. This capability can be harnessed and lets a sound card be utilized as a low-frequency waveform generator. In most cases, no hardware modifications are needed and arbitrary electrical signals can be readily generated, provided that the sound card is driven with the right software components. Naturally, a sound card based waveform generator is not as capable as a dedicated hardware solution and has many limitations such as a narrow range for the generated voltage. However, the nature of ECG signals, which are inherently low-frequency and on the order of a few hundred microvolts, can be generated by a sound card without any problems The majority of dedicated sound cards as well as devices with integrated sound support have output frequencies of up to 20kHz, which is well than enough. A software waveform generator based on a sound card is therefore a viable option for signal injection. It not only drastically reduces cost, but also simplifies the injection method. Figure 5 depicts a possible setup where a software waveform generator is run on a laptop that injects the generated signal though its audio output port.

3) Audio Playback: Instead of using a software waveform generator and changing the function of the sound card, we explore the possibility of playing back stored ECG signals on the sound card as actual sound. Such an approach does not require specialized software, i.e., a software waveform generator, and might be executed on any device capable of outputting analog audio signals. This could reduce effort and complexity of a presentation attack to a great extent.

The challenge of replicating an ECG signal directly as audio output consists of transforming the digital representation of an ECG signal into an audio file that can be played back on the sound card. We wrote software that filters the ECG signal, applies the correct scaling of voltage levels, sets the sampling rate and finally stores the signal as an audio file (WAV format).
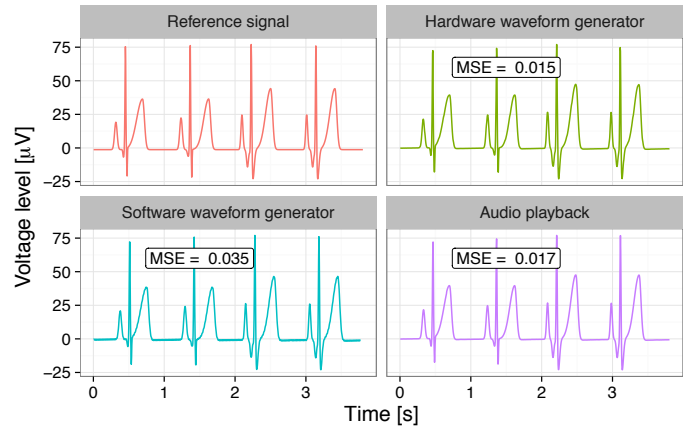

Fig. 6: Reference ECG signal compared to the ECG traces measured when the reference signal is injected using three different injection methods. The traces are captured by the Nymi Band and read out with the Nymi software development kit.

The resulting file can then be played back on almost any device and potentially injected into the sensors of a biometric system based on ECG recognition.

The contraption we used for the evaluation of the audio playback as injection method is identical to the hardware setup in Figure 5. Nevertheless, the attack can be carried out with any device capable of analog audio output.

C. Injection Quality

In order to validate the presented signal injection methods and assess their quality, we select a stored reference signal, reproduce and inject it using each of the three approaches. We then compare the reference signal to the traces the electrocardiograph measures while injection takes place. In case of the Nymi Band, the captured traces can be accessed

	Approximate cost	Required knowledge	Physical size	Signal quality
– Hardware waveform generator				
	$240	high	large	very high
– Software waveform generator				
	$50	moderate	small	high
– Audio playback				
	$10	moderate	very small	high

TABLE I: Comparison of injection methods

and read out with the Nymi software development kit (SDK).

In addition to a visual comparison between the stored signal and the extracted ECG traces, we verify their similarity numerically. The distance metric for comparison is the (per sample) mean squared error (MSE). We make sure the reference signal's sampling rate matches that of the Nymi Band. Also, the measured traces might be linearly translated and require alignment before the calculation of the distance metric. We determine the constant shift between the stored reference signal and captured traces by aligning the peaks of the R waves.

Results are show in Figure 6. All three proposed injection methods manage to reproduce the reference signal remarkably well and we conclude that the contraptions are effective. It is evident that signal quality is proportional to sophistication and cost of the contraption used for signal injection: The hardware waveform generator achieves the smallest error (mean squared error of 0.015) as opposed to the audio playback (which is slightly less accurate and achieves an MSE of 0.017) and the software waveform generator.

D. Comparison of Injection Methods

In Table I, we present a comparison between our three injection methods along the criteria outlined above (see Section III-A). Not surprisingly, the hardware waveform generator achieves the highest signal quality, but at the same time entails the highest cost. Entry-level arbitrary waveform generators retail at around $250. They are, however, fairly bulky and only designed for stationary use.

Software waveform generators are not only available for personal computers, but even smart phones. A low-end smart phone equipped with an analog audio output costs around $50 as of spring 2016. An ECG signal encoded as an audio file can even be played back on a cheap portable audio player which can cost less than $10 and has a tiny form factor.

IV. EXPERIMENTAL DESIGN

In this Section we will outline a number of approaches that can be used by an attacker to obtain data for a presentation attack. Based on these attack vectors, we will then discuss our data collection methodology.

A. Obtaining Data for a Presentation Attack

In Section III we have demonstrated our capability to inject arbitrary signals into the Nymi Band. However, the attacker still requires an input ECG signal that is sufficiently close to


Fig. 7: The ECG monitor data is extracted through image analysis of this plot shown by the software. This models the threat of an attacker obtaining a photo of the victim's ECG.

that of the victim. There are multiple conceivable approaches to obtain this data:

Medical records often contain printouts of a patient's ECG. A conventional hospital ECG uses 10 adhesive electrodes to simultaneously record 12 leads (see Section II for details). An attacker could obtain these records either electronically (e.g., through social engineering or a compromised medical database) or on paper (e.g., by taking a photo of the plots). In addition, mobile devices that allow patients to monitor their health at home are becoming more widespread. Besides ECG monitors for medical use, a number of devices are marketed for fitness, for example in the form of heart rate monitors used during cardiovascular exercise. In these cases, the data could be intercepted during transmission (e.g., to the victim's smartphone) or leaked through an insecure or malicious mobile app.

The Nymi SDK allows NEAs to collect arbitrary amounts of raw ECG data (see Section II). There are two conceivable ways in which this might pose a security risk: NEAs have to be actively paired with the NCA by the owner of the band. However, this only means they are trusted by the user, not that they are inherently trustworthy. A rogue NEA could trick the user into providing (a sample of) his ECG, which would then allow the owner or developer of the NEA to later use this data to carry out a presentation attack. A second, probably more severe, way of abusing the SDK data collection is through a social engineering attack. Given the novelty of the Nymi Band an attacker might ask the victim to wear the attacker's Nymi Band to test whether it is possible for the victim to unlock it. Instead of performing regular activation, the attacker could instead activate the band through the backup password and then collect the victim's ECG data through a previously paired NEA (e.g., on a laptop or the attacker's smartphone). This attack is particularly dangerous, as the data is collected directly on a Nymi Band, rather than a different device that requires the application of the mapping function (see Section V).

B. Data Collection

The previous subsection has outlined several scenarios that might allow an attacker to obtain a victim's ECG data. We collect data using a variety of devices to reflect them:

The first device we use is a lightweight medical ECG monitor, the Heal Force Prince 180B (see Figure 8). The device has two main measurement modes which use either the built-in electrodes on the sides of the device or an external 3-lead ECG cable with disposable electrodes. The first measurement (which we will refer to as the Palm measurement throughout the paper) uses the built-in electrodes. Participants were asked to hold the device as pictured in Figure 8 and to remain still during the measurement as the ECG recording is highly sensitive to device movements. When using the built-in electrodes, the device always records data for a fixed duration of 30 seconds. Following the palm measurement, we use the 3-lead cable to record Lead I and Lead II (see Section II for details), which involves attaching the disposable electrodes to both arms and the left leg. Unlike hospital ECGs, which capture all 12 leads simultaneously, this monitor is limited to recording a single lead. However, by switching the position of the electrodes, all standard leads can be recorded in sequence. We chose to record Lead I, which has measurement points similar to the Nymi Band, and Lead II, which measures the potential difference between the right arm and left leg. Due to practical reasons we choose not to collect a full 12-lead ECG. Both leads are recorded during a medical ECG, as such an attacker could obtain them by taking a photo of the patient's medical files. In order to reflect this, we don't extract raw data from the device, but instead obtain it by performing image analysis on the plots displayed by the software (see Figure 7).

The second device, pictured in Figure 9, is a mobile ECG monitor that can be attached to a smartphone and transmits the recorded data via Bluetooth. Following the measurement, the device provides an instant assessment with regard to a number of heart disorders. Participants were again asked to remain still during the 30-second measurement period to avoid the introduction of additional noise. Following the measurement, the device creates a pdf report which can be automatically sent to an arbitrary e-mail address (such as the patient's physician). This report contains, aside from the patient's personal information, a plot of the recorded ECG data. The report is sent via e-mail without any encryption or other security features. Similar to the ECG monitor data, we use image analysis to extract the raw data from the pdf file.

Lastly, we collect data using the Nymi Band itself. As outlined in Section II, the Nymi SDK allows Nymi Enabled Applications (NEAs) to collect raw ECG data once the Nymi Band is unlocked and paired with the NEA. As a result of the band's hardware design, this requires the cooperation of the user as they have to touch the top electrode of the band to enable ECG recording. Following the user enrollment and activation of the Nymi Band we collect 60 seconds of raw ECG data using the SDK. We used three (identical) developer kits of the Nymi Band running the SDK version 1.03. It is important to note that the capability of the band to report raw ECG data to NEAs has been discontinued from SDK version 2.0 onwards. However, it is still possible to collect data by using the legacy SDK on the developer bands. The final consumer version of the band has been released in September 2016 and also lacks the

Fig. 8: ECG monitor in palm measurement mode

Fig. 9: Mobile ECG monitor

capability to record raw ECG data. We can't make any claims on the success of our attack on this new version, although, based on the published changes in the consumer version, we can not see any structural obstacles to the attack still being successful.

C. Participant Recruitment and Ethical Considerations

This project has been reviewed by and received clearance from the Central University Research Ethics Committee of the University of XXXX, reference number YYYY[2]. The main ethical concern when gathering ECG data is the sensitive nature of the data itself. This sensitivity stems from the fact that a variety of heart disorders can be diagnosed using ECG, including disorders the participant may not have previously been aware of. The possibility of false positives (i.e., the incorrect diagnosis of a condition) in conjunction with the fact that none of the researchers are trained medical professionals led us to the decision of disabling all diagnostic capabilities of the devices used and inform all participants accordingly. Since a future diagnosis based purely on the data is theoretically possible, we store all datasets anonymized to erase any links between a potential condition and a single participant. Given the above concerns, we required participants to be at least 18 years old as the only criteria for inclusion in the study.

We recruited a total of 41 participants (21 female, 20 male) through mailing list adverts and social media. Participants were

---

[2]Blanked during peer review

made aware that the research involves ECG biometrics, but were only told the specific purpose of the data collection afterwards. Participants were paid in cash for their participation.

V. MAPPING FUNCTION

In this section, we show how to generate attack signals for the Nymi Band using ECG data from different sources. The method is based on the derivation of *mapping functions*, i.e., functions that transform signals recorded from one device, the *source*, in order to resemble the morphology of the signals from a *target* device. In our case, the target device corresponds to the Nymi Band. In other words, we aim to find a function that, given in input an ECG signal from a source device, is able to produce the "same" signal as if it was recorded on the target device.

In this way, the mapping function can mitigate and even eliminate the statistical differences in biometric features that arise from using different measurement devices and modes. In Figure 11, we compare the mean ECG signals recorded for the same individuals but from different devices, showing that the signals exhibit device-specific morphologies, especially as far as amplitude features are concerned. For instance, we observe that for both subjects the palm measurements yield a more prominent T wave, while Lead II signals yields the lowest T wave peak. Similarly, in both cases the P waves obtained from Lead I and the mobile ECG monitor stand above those of the Nymi Band and Lead II, and the R wave of Lead II has the least amplitude. These observations demonstrate that many discrepancies in the ECGs are device-specific and thus can be addressed by the application of mapping functions.

Figure 10 shows a summary of the methods for estimating mapping functions and for generating attack signals. Let S and T be the source and target devices, respectively. Let J be the set of ECG features described in Figure 1, and I be the set of subjects we use for computing the mapping. The training dataset consists of the sets $\{\mathsf{ECG}_i^S\}_{i \in I}$ and $\{\mathsf{ECG}_i^T\}_{i \in I}$ of ECG signals recorded, for each subject $i \in I$, with the source and target device, respectively. The method is based on the following steps:

1) **Feature extraction.** From the input ECG data, we extract the relevant biometric features. The outputs of this step are, for each subject $i \in I$, sets of discrete probability distributions $\mathcal{D}_i^S = \{D_{i,j}^S\}_{j \in J}$ and $\mathcal{D}_i^T = \{D_{i,j}^T\}_{j \in J}$, where $D_{i,j}^S$ $(D_{i,j}^T)$ is the distribution of ECG feature j for subject i in the source (target) signal. Specifically, we consider the time domain features summarized in Figure 1 and apply the algorithm of [2] for their detection.

2) **Mapping estimation.** This boils down to an optimisation problem (described in Section V-A) where we seek to find an optimal mapping, i.e., a set of transformation functions $\mathbf{f} = \{f_j\}_{j \in J}$ with $f_j : \mathbb{R} \to \mathbb{R}$, such that, for each feature j and subject i, they minimise the statistical distance between the transformed source distribution $f_j(D_{i,j}^S)$ and the corresponding target distribution $D_{i,j}^T$ [3]. In other words, f_j transforms values

---

[3]Technically, for discrete distribution D, $f_j(D)$ is the distribution whose support is the image of $\mathsf{supp}(D)$ under f_j ($\mathsf{supp}\,(f_j(D)) = f_j[\mathsf{supp}(D)]$) and with probability mass function defined, for $x \in \mathsf{supp}\,(f_j(D))$, by $f_j(D)(x) = \sum_{x' \in f_j^{-1}[x]} D(x')$ where $f_j^{-1}[x]$ is the preimage of x under f_j, that is, all the elements $x' \in \mathsf{supp}(D)$ such that $x = f_j(x')$.

of feature j from device S in order to be as close as possible, statistically speaking, to the values of the same feature from device T. We restrict the search to linear functions, of the form:

$$f_j(x) = a_j x + b_j \qquad (1)$$

Note that linear mappings are adequate in this context because the amplitudes of the ECG wave peaks along different leads are linearly related [14]. Moreover, unlike more complicated transformation functions (e.g. polynomial or logarithmic), linear mappings do not suffer from over-fitting problems when using small training datasets [11].

Figure 12 compares the distributions of a selection of ECG features before and after applying the mapping function. In this case, Lead II is the source device and the Nymi Band is the target. Input data is obtained from a subject in our training set. The plots demonstrate that, after the transformation, the distributions of source features (blue bars) practically overlap with the corresponding target distributions (red bars), consistently reducing the statistical differences observed for the initial, non-transformed, source features.

Once estimated, the mapping $\mathbf{f}$ between S and T is used to generate attack signals for device T starting from new signals recorded with S. Let $i' \notin I$ be our victim, for which we possess an S-signal $\mathsf{ECG}_{i'}^S$. The procedure, illustrated in Figure 10 (b), consists of the following steps:

1) Extract the feature distributions $\mathcal{D}_{i'}^S$ from the source signal $\mathsf{ECG}_{i'}^S$.

2) Apply the estimated mapping function $\mathbf{f}$ to derive the transformed features distributions: $\mathbf{f}(\mathcal{D}_{i'}^S) = \{f_j(D_{i',j}^S)\}_{j \in J}$.

3) Produce an attack signal by generating a synthetic ECG signal out of the transformed features $\mathbf{f}(\mathcal{D}_{i'}^S)$, as explained in Section V-B.

In Figure 13, we compare the original source signal, the corresponding attack ECG generated from the transformed features and the target signal of a participant from the training set. We observe that the attack signal prominently improves on the original one, in some sections being virtually identical to the target signal that we aim to reproduce. This demonstrates the effectiveness of our synthetic ECG generation: as the transformed ECG features get closer to the target features, the generated attack signal better emulates the target signal. Nevertheless, the mapping function is designed to be as general as possible and to be effective for the whole training dataset, which explains why not all subject-specific differences can be eliminated, see for instance the higher P wave of the attack signal in plot (a), or the lower T wave in plot (b).

We would like to stress the generality of our method, which mostly relies on estimating linear transformations between sets of biometric features. Indeed, the method can be easily extended to support other classes of biometrics, provided the availability of algorithms for feature extraction and generation of synthetic signals.

A. Optimization Problem

We formulate the problem of finding the best mapping function as a non-linear constrained single-objective optimiza-
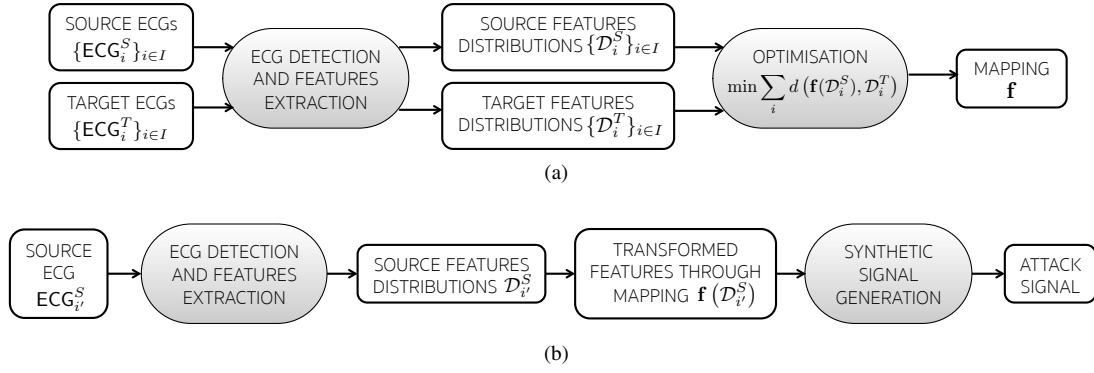
Fig. 10: Methods for the derivation of optimal mappings (a) and generation of attack signals (b).
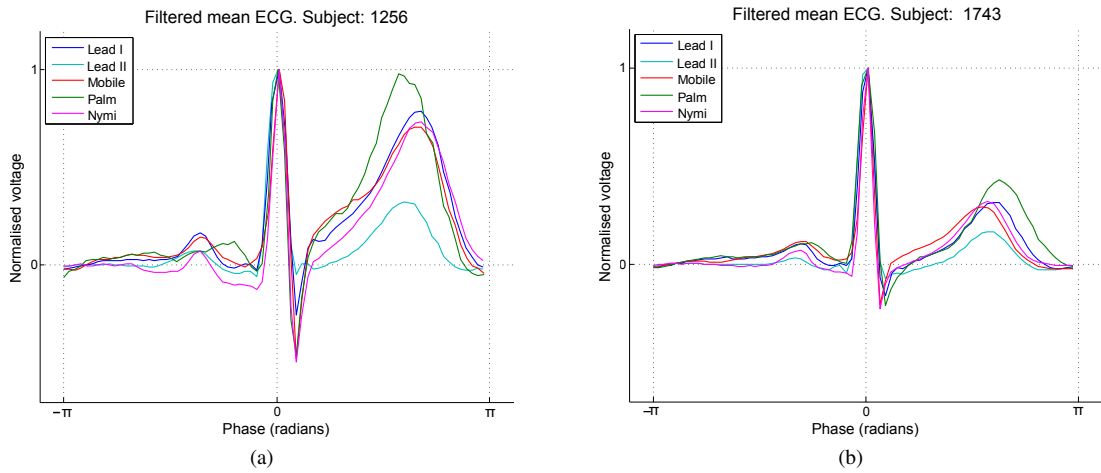
Fig. 11: Comparison of mean ECGs from same subjects and different leads/devices. Mean ECGs are computed after a linear phase assignment [19], assigning a periodic phase value to each sample in the ECG, starting from one R-peak (phase 0) and ending with the next R-peak (phase 2π). For each heart cycle, amplitudes are scaled by the amplitude of the corresponding R-peak.

Fig. 12: Feature distributions before (first row) and after (second row) the application of the mapping function. T width (L) and (R) denote the left and right component of the T width (see Figure 1).

(a) Mobile	(b) Palm

Fig. 13: Comparison of mean ECGs among source signal, corresponding attack signal and target Nymi signal.

tion problem that we solve using a genetic algorithm [9]. The problem is defined as follows:

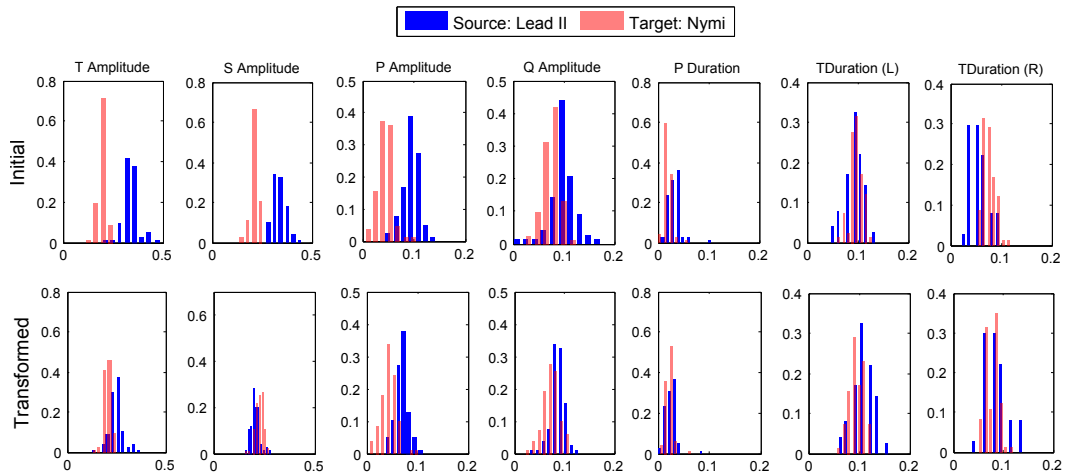$$\underset{(a_j,b_j)_{j\in J}}{\text{minimize}} \quad \sum_{i\in I^*} d\left(\mathbf{f}(\mathcal{D}_i^S), \mathcal{D}_i^T\right) \qquad (2)$$

$$\underset{i\in I, j\in J}{\text{subject to}} \quad a_j, b_j \in [k_j^\perp, k_j^\top] \qquad (3)$$

$$a_j \cdot D_{i,j}^{S,\min} + b_j \in \left[D_{i,j}^{T,\min *}, D_j^{T,\max *}\right] \qquad (4)$$

$$a_j \cdot D_{i,j}^{S,\max} + b_j \in \left[D_{i,j}^{T,\min *}, D_j^{T,\max *}\right] \qquad (5)$$

The decision variables are, for each feature $j \in J$, the linear coefficients a_j and b_j characterising the transformation f_j we seek to estimate (see Equation 1). The objective function is the sum over subjects $i \in I^* \subseteq I$ of the distance between the transformed source distributions of i, $\mathbf{f}(\mathcal{D}_i^S)$, and the corresponding target distributions $\mathcal{D}_i^T$. Here d is a generic distance measure (discussed later). In particular, I^* is a subset of the training set I and is obtained as follows: 1) we compute the distances $d\left(\mathbf{f}(\mathcal{D}_i^S), \mathcal{D}_i^T\right)$ for all $i \in I$; 2) we apply the Grubbs' test [10] to detect the set of outliers $I' \subseteq I$ on these distances; 3) we remove the identified outliers: $I^* = I \setminus I'$. These three steps are repeated until a maximum number of outliers is removed or no further outliers are identified. The rationale for considering I^* instead of I in the objective function is that we do not want to penalise mappings that perform well for most subjects and poorly for few of them (the outliers). This approach has the added advantage to bypass subjects with inaccurate input ECG data, e.g., through noise introduced through excessive movement. These cases are indeed very likely to be identified as outliers.

Regarding the feasible region of the optimization problem, Equation 3 ensures that the linear coefficients are bounded in some real-valued interval $[k_j^\perp, k_j^\top]$. The purpose of Equations 4 and 5 is to constrain the ranges of the transformed source features, in a way that they are similar to the ranges of the target features. Preliminary results showed that these constraints are crucial to ensure that the corresponding attack signal resembles a biologically realistic ECG. For subject i and feature j, let

$D_{i,j}^{T,\min}$ and $D_{i,j}^{T,\max}$ be the minimum and the maximum values of distribution $D_{i,j}^T$, respectively[4]. We define the lower and the upper bounds for the transformed features as:

$$D_j^{T,\min *} = (1-q)\cdot\min_{i\in I} D_{i,j}^{T,\min} \quad \text{and} \quad D_j^{T,\max *} = (1+q)\cdot\max_{i\in I} D_{i,j}^{T,\max}$$

where $q \in (0,1)$ is a factor for relaxing the range width. The resulting range constraints are given, for all points x in the support of the source distribution $D_{i,j}^S$, by:

$$a_j \cdot x + b_j \in \left[D_j^{T,\min *}, D_j^{T,\max *}\right]. \qquad (6)$$

Note that the number of such constraints quickly explodes with the number of subjects, features and distinct data points per feature. However, by the monotonicity of the linear mappings, it suffices to check Equation 6 only for the minimum and maximum values of $D_{i,j}^S$, denoted respectively by $D_{i,j}^{S,\min}$ and $D_{i,j}^{S,\max}$, thus yielding Equations 4 and 5. Importantly, this implies that our estimation method supports not just linear functions, but general monotonic functions.

Statistical distance. The distance function of Equation 2 is defined as the mean of the statistical distances between the transformed and the target distributions over all the features:

$$d\left(\mathbf{f}(\mathcal{D}_i^S), \mathcal{D}_i^T\right) = \frac{1}{|J|} \sum_{j\in J} d_s\left(f_j(D_{i,j}^S), D_{i,j}^T\right).$$

where d_s is a generic statistical distance. Among the possible candidates for d_s, we chose the L^2 distance between distributions. Let $\mathcal{F}_{i,j}^{\overline{S}}$ and $\mathcal{F}_{i,j}^T$ be the piece-wise linear estimations of the cumulative distribution functions of $\mathbf{f}(D_{i,j}^S)$ and $D_{i,j}^T$, respectively. Then, we define d_s as the L^2 distance between functions $\mathcal{F}_{i,j}^{\overline{S}}$ and $\mathcal{F}_{i,j}^T$ [22, Chapter 1]:

$$d_s\left(f_j(D_{i,j}^S), D_{i,j}^T\right) = w_j\left(\int_{D_j^{T,\min *}}^{D_j^{T,\max *}} \left(\mathcal{F}_{i,j}^{\overline{S}}(x) - \mathcal{F}_{i,j}^T(x)\right)^2 dx\right)^{\frac{1}{2}}$$

where $w_j = D_j^{T,\max *} - D_j^{T,\min *}$ is introduced as a normalisation factor. In the implementation, the above integral is approximated using a composite mid-point quadrature formula.

---

[4]With abuse of notation, the minimum and maximum of a discrete distribution D are meant as the minimum and maximum of its support.

B. Synthetic Signal Generation

Synthetic ECG signals are generated as the sum of Gaussian functions, used to reproduce the typical bell-shaped curves of the ECG waves and parametrised by sampling values from a given set of feature distributions. As previously explained, for attack signals we consider ECG features after the application of some mapping function.

The method extends [16], [2] in order to support asymmetric ECG waves, which are physiologically more accurate, thus leading to attack signals that better emulate the Nymi Band's ECG, as discussed in Section V-C.

Let $(\mathsf{PP}_1, \ldots, \mathsf{PP}_{n-1})$ be the sequence of PP intervals detected from the source signal. The sequence is used to determine the beginning of each heart cycle such that, for $h = 1, \ldots, n$, the h-th heart cycle starts at time $T_h = T_0 + \sum_{k<h} \mathsf{PP}_k$, where T_0 is the offset of the first P wave.

For each ECG wave kind $w \in \{P, Q, R, S, T\}$ and heart cycle $h = 1, \ldots, n$, the considered features are: wave amplitude, $A_{w,h}$; wave peak location relative to the start of the h-th cycle, $L_{w,h}$; and left and right components of the wave width at half amplitude, $W_{w,h}^l$ and $W_{w,h}^r$. Note that peak locations are easily derived from the interval features shown in Figure 1. The synthetic ECG at time t is defined as follows:

$$s(t) = \sum_{h=1}^{n} \sum_{w \in \{P,Q,R,S,T\}} G\left(t, T_h + L_{w,h_1}, A_{w,h_2}, W_{w,h_3}^l, W_{w,h_4}^r\right)$$
(7)

where $G\left(t, T_h + L_{w,h_1}, A_{w,h_2}, W_{w,h_3}^l, W_{w,h_4}^r\right)$ is the value at point t of an asymmetric Gaussian curve centred at $T_h + L_{w,h_1}$, with amplitude A_{w,h_2}, and full width at half maximum made of left component W_{w,h_3}^l and right component W_{w,h_4}^r. G is given by:

$$G\left(t, L, A, W^l, W^r\right) = A \cdot \exp\left(-4 \cdot \log 2 \cdot \frac{(t-L)^2}{W(t)^2}\right)$$

where $W(t) = W^l$ if $t \leq L$ and $W(t) = W^r$ otherwise.

Note that in Equation 7 the features used to generate the Gaussian curve are not necessarily drawn from the same heart cycle h. Specifically, for each cycle h, we randomly sample the heart cycles $h_1, \ldots, h_4$ from which location, amplitude and width features are extracted. Based on preliminary results, among the possible sampling strategies, we choose peak location and widths from the same heart cycle, i.e., $h_1 = h_3 = h_4$.

Importantly, such generated synthetic ECGs account for the specific inter-beat patterns of the subject (another common ECG biometric feature), since we use the same PP sequences detected from the source signal.

C. Evaluation

In this section, we perform an in-depth evaluation of the methods for estimating mapping functions and generating synthetic signals. The aim of the following experiments is to obtain insight into HeartID, the Nymi Band's authentication and biometric recognition library, in order to devise the best design choices for our methods, such as the ECG features to include in the mapping or the filtering algorithm to use in the ECG detection procedure. There is only very little information available about the algorithms used in HeartID apart from [6] which proposes a continuous biometric recognition system based on ECG called HeartID.

Unfortunately, it remains unknown to what degree the techniques and algorithms described in [6] have been included in the authentication library currently being used by the Nymi Band's companion app. As such, we do not have any prior knowledge of the classifiers or features used in the authentication process. Nevertheless, we can use the NCA as an oracle by querying it with an ECG signal and recording the response (i.e., an accept or a reject of the signal).

The main obstacle to such an extensive analysis is the time needed to inject attack signals using a waveform generator or a sound player. To overcome this limitation we devised another kind of attack, called *offline attack*, which is instantiated by directly interfacing with HeartID through the API calls of the NCA (Nymi Band's companion app). We implemented a simple Android app that allows setting up previously stored biometric templates and performing authentication for arbitrary attack signals, at a rate of hundreds of signals per minute, without requiring a waveform generator or Nymi Band. Our devised Android app does not require physical ECG input, but accepts biometric template and biometric samples in digital form via a command line interface and forwards them to the HeartID library. For every authentication attempt, the library's authentication decision is stored in a database by our Android app for later analysis.

We remark, however, that the final results of Section VI are obtained online, i.e., using the actual Nymi Band.

The experiments below were performed on a selection of 8 subjects from the training set. For each experiment, subject and device, we tested 20 randomly generated attack signals. We recall that the synthetic ECG is generated by sampling features from distributions, hence the reason of their randomness. The reported success rate (SR) is the ratio between successful authentications and total attempts. As expected, the enrolment signals used to build the biometric templates yielded an SR of 100%.

Filtering. We evaluate the adequacy of our ECG filtering algorithm, which builds on a Savitzky-Golay (SG) smoothing filter [17]. We applied the SG filter to the enrolment signals, obtaining for them an SR of 100%. Stronger SG smoothing parameters resulted in SRs below 63%. We also wanted to assess how noise affects authentication chances. To this purpose, we added to the filtered signal white Gaussian noise with standard deviation computed from the enrolment signal. These yielded an SR of 100%, demonstrating that the filtering method of HeartID can equally support noisy and filtered ECG data. Note that the above experiments did not require the generation of synthetic ECGs.

Synthetic signal generation. Here, we assess the fit of our synthetic ECG signals. To this purpose, we generated signals drawing on the features detected in the enrolment ECGs, and reproduced in the same order. These signals resulted in an overall SR of 75%, with an SR of 100% for 5 out of 8 subjects. To understand if HeartID is sensitive to all the ECG waves, we further produce signals where one of the waves is systematically suppressed by setting its amplitude to 0. For each wave kind, the obtained SR was 0%, suggesting that all ECG waves are used

in the biometric template. In a variation of the first experiment, we test signals with symmetric waves, resulting in an SR drop from 75% to 15%. This motivates our claim that asymmetric waves are needed in order to produce realistic and successful attack signals. Finally, we report that subject-specific ECG features are indeed central for authentication: signals generated using default parameters from literature [4] produced an SR of 0%.

Statistical distance. Now that we have proved the adequacy of the filtering and synthetic signal generation algorithms, the next step is evaluating the mapping function. In this experiment we want to assess how different statistical distances in the optimisation problem of Section V-A affect the success rate. Here, we restrict to *one-2-one mappings*, i.e., mappings estimated over data from a single subject. In contrast to the default mappings, one-2-one mappings are tailored to the subject, meaning that the resulting attack signal better mimics the target ECG. However, these cannot be used to instantiate an attack, since they require the victim's enrolment signal. We compare the previously introduced L^2 distance with the χ^2 index for empirical distributions and the Kolmogorov-Smirnov statistic (KS) [5]. Using L^2, we get an overall SR of 49.9%, while for χ^2 and KS, the SR is 29.5% and 39.6%, respectively. This supports our choice of L^2.

Mapping strategies. We finally evaluate a number of alternatives to the default mapping functions. Signals generated using the default settings, where the mapping is estimated from the whole training dataset and considers the full set of ECG features, resulted in an SR of 26.2%. Slightly worse success rates are obtained when some of the features are not transformed, i.e., left as detected in the source signal: the overall SR is 22.8%, 21.1% and 20% if excluding, respectively, widths, peak intervals and both. As expected, the exclusion of amplitude features has shown no significant success. We get comparable results when considering sub-optimal mappings in place of the best mapping function found by the optimisation algorithm: for the second and third best mappings, we have SRs of 21.5% and 21.9%, respectively. We also learnt that one-2-one mappings estimated from one single subject are mostly ineffective for attacking a different subject (SR = 4.4%), thus confirming the importance of gathering a substantial training dataset. Another alternative that performed poorly (SR = 11.5%) is building a mapping from a single "super-individual", obtained by merging the features distributions across all subjects.

VI. RESULTS

We conduct the signal injection attack by combining the building blocks described in the previous sections. After enrolling a user into the Nymi Band, we first inject the raw data that was collected using each of the devices (see Section IV). Following that, we apply the mapping function, which has been trained on the initial set of users (i.e., not including the user that is being attacked) to obtain the transformed signals. As outlined in Section V, the mapping function employs random sampling of features to generate the attack signal. We use between 2 and 4 repetitions of this sampling process and report whether at least one of them was accepted by the Nymi Band. As the Nymi Band does not limit the number of authentication

User	Lead I	Lead II	Mobile	Palm	Combined	Nymi
1	✗	✗	✗	✗	✗	✓
2	✗	✓	✗	✗	✓	✓
3	✗	✗	✗	✗	✗	✗
4	✗	✗	✗	✗	✗	✓
5	✗	✗	✓	✗	✓	✓
6	✗	✗	✗	✗	✗	✓
7	✗	✗	✗	✗	✗	✓
8	N/A	N/A	✓	✗	✓	✗
9	N/A	N/A	✗	✗	✗	✓
10	✗	✗	✗	✗	✗	✓
11	✗	✗	✗	✗	✗	✓
12	✓	✗	✓	✗	✓	✓
13	✓	✗	✗	✗	✓	✓
14	✗	✗	✓	✗	✓	✓
15	✓	✗	✗	✗	✓	✗
16	✗	✗	✗	✗	✗	✗
17	N/A	N/A	✗	✗	✗	✓
18	✗	✗	✗	✗	✗	✓
19	✗	✗	✗	✗	✗	✓
20	✓	✗	✓	✗	✓	✓
21	✓	✗	✓	✗	✓	✓

TABLE II: Results of injecting the unmodified raw data. "Combined" refers to the combination of all non-Nymi devices. The tick marks signify users that were successfully attacked.

User	Lead I	Lead II	Mobile	Palm	Combined
1	✗	✗	✗	✓	✓
2	✗	✓	✗	✗	✓
3	✗	✗	✗	✗	✗
4	✗	✗	✗	✗	✗
5	✓	✗	✓	✓	✓
6	✗	✗	✗	✗	✗
7	✗	✗	✗	✗	✗
8	N/A	N/A	✓	✗	✓
9	N/A	N/A	✗	✗	✗
10	✗	✗	✓	✗	✓
11	✗	✗	✗	✗	✗
12	✓	✗	✓	✗	✓
13	✓	✗	✗	✗	✓
14	✓	✗	✗	✗	✓
15	✓	✗	✓	✓	✓
16	✗	✗	✗	✗	✗
17	N/A	N/A	✗	✗	✗
18	✓	✗	✗	✗	✓
19	✓	✗	✗	✓	✓
20	✓	✗	✓	✗	✓
21	✓	✗	✓	✓	✓

TABLE III: Results of injecting the data generated by applying the mapping function. The shaded cells show the results that differ from Table II. Due to source and target device being identical, the mapping function is not applied to the Nymi SDK data. The tick marks signify users that were successfully attacked.

attempts a user can make, the only cost of repetitions lies in an increased time required to carry out the attack. Due to space constraints we report only the results of injecting the ECG signals by playing back .wav files using an off-the-shelf music player. As we have discussed in Section III, the performance that can be expected is similar for all three devices, and the .wav playback poses the lowest technological barrier for the attacker, both in terms of cost and compactness. The raw results of our analysis are shown in Table II and III, a summary is shown in Figure 14. For 3 out of 21 users we were unable to collect high-quality signals using the external electrodes of the ECG monitor, leaving us with 18 attack attempts for the Lead I and Lead II data sources.

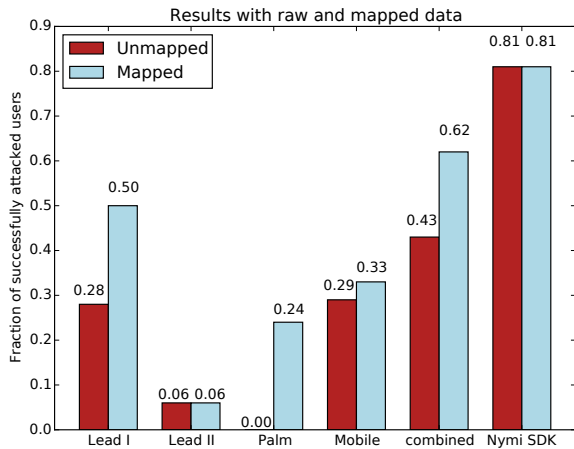Not surprisingly, the success rate of injecting data collected

Fig. 14: Fraction of users successfully attacked. Results are shown for raw data and the data obtained by applying the mapping function.

with the Nymi Band is the highest. With the exception of four users, using this data resulted in unlocking the band, leading to a success rate of 81%. However, two out of the four unsuccessful users (3 and 6) failed to authenticate themselves following enrolment. This hints that the failure of the attack is most likely a consequence of erratic or noisy data being collected during the enrolment phase. Injecting the Lead II measurement succeeded only for a single user. Intuitively, this makes sense since Lead II measures voltage between the left leg and right arm (rather than the left and right arm, which is approximately what the Nymi Band observes). Conversely, Lead I performs relatively well in the attack, succeeding in five out of 21 users (24%), most likely being a result of the similar measurement points. Based on this intuition, one would expect the palm measurements to perform similarly well, as the measurement points (palm of the left hand and index finger of the right hand) are even closer to those used by the Nymi Band (wrist of left hand and index finger of right hand). However, using these measurements caused the attack to fail for all users (and is the only data source to do so).

Applying the mapping function considerably improves these results, particularly for those data sources that performed poorly initially. The success rate for using Lead I data improves from 28% to 50%. The Palm measurements, which were initially unsuccessful for all users, could be used for successful attacks on 5 users, thus increasing the success rate from 0% to 24%. The mobile ECG monitor data source is the only one where applying the mapping function causes the attack to fail for one user. This is most readily explained by the unmodified data just being on the edge of being accepted. The only data source for which the mapping function had no effect (positive or negative) is the Lead II data obtained from the ECG monitor. This could be either due to limitations of the mapping function (such as the precise feature set of the Nymi Band being unknown), or due to important biometric information simply not being present in this lead. The former case could possibly be remedied by obtaining a better understanding of the biometric features involved (which is difficult in a blackbox scenario as presented

by the Nymi Band). In the latter case, it is not possible to find a mapping between the feature distributions, thus requiring the attacker to obtain a different source of ECG information. In the medical domain it is necessary to measure multiple ECG leads as they contain different kinds of information, so it is not implausible that this is similar for identifying (biometric) information.

Overall, the attacker's chance of success is 81% assuming they have obtained a measurement through the Nymi Band, and 62% if they have only obtained data from the remaining sources. The latter is computed as the fraction of users for which at least one of the four data sources led to a successful activation of the band (shown as the "Combined" column in Table II and III)

VII. RELATED WORK

The increasing popularity of biometrics as an authentication mechanism has sparked research into a variety of attacks. Multiple approaches to imitate users or forge measurements have been proposed, including attackers modifying their own behavior and attacks being carried out automatically.

Sharif et al. demonstrate an attack on face recognition systems that relies on the attacker wearing a pair of printed eyeglass frames [21]. Their approach enables the attacker to both impersonate a victim and to evade face detection. The attack treats the classifier as a whitebox (i.e., assumes knowledge of the internals) and highlights the danger of using machine learning in adversarial settings.

Tey et al. propose an attack on keystroke dynamics (distinctive typing patterns), targeting systems using the biometric for password hardening [25]. The attack assumes the password to be known to the attacker, who then has to mimic the legitimate user's keystroke dynamics while typing it. The authors devise a system that gives positive and negative feedback to the attacker regarding their closeness to being accepted. This feedback relates to individual features, such as character hold times and interkey times (i.e., the time between pressing two individual keys). Depending on the information given to the attackers, a false accept rate of up to 99% is achieved, suggesting that it is possible for humans to adapt their own typing patterns enough to fool state-of-the-art systems. Nevertheless, this approach assumes that both the password and the typing patterns are known, which requires either use of a keylogging device or a password database compromise. In addition, it is doubtful whether this approach can be adapted to continuous authentication systems that work on free texts, rather than short passwords.

Touchscreen input biometrics have gained immense popularity over the past few years due to the rise of smartphones, tablets and other touchscreen devices. Zheng et al propose an authentication system based on touchscreen input and also measure the advantage gained by attackers through observing the victim [27]. The results indicate that simple observation is not enough to trick the system, although their work assumes that the attackers have no specific knowledge of the system's features. As such, the failed impersonation could be due to insufficient information, or the fact that touchscreen inputs are inherently difficult for humans to mimic. Serwadda et al. take a different approach by building a Lego robot to carry out the

attack, rather than having a human modify their behavior [20]. Instead of using biometric data from the targeted legitimate user, they create a "catch-all" attack based on population statistics. Using a subset of all users, they train the robot to perform swipes and scrolls in a way that is sufficiently common for a high number of users. Carrying out the attack increase the system's Equal Error Rate (EER) by between 339% and 1004%. At present, authentication systems based on touchscreen inputs are not commercially available, causing the authors to create their own. While they loosely base their features and system design on the popular work by Frank et al. [7], the initial (i.e., before the attack) EER (13%) of the system is much higher even after excluding users exhibiting poor error rates. Unlike a regular single-target imitation attack, the proposed attack targets multiple users at once and therefore requires an overlap between users' templates that would be unlikely to be found in systems with lower error rates. As such it is doubtful whether the described attack is as effective or applicable at all when using state-of-the-art systems.

VIII. Discussion and Countermeasures

In the previous sections we have outlined an effective presentation attack against ECG biometrics in general and have shown its effectiveness when applied to the Nymi Band. Most generally, there are two main approaches to mitigate the attack:

The first approach is liveness detection. Liveness detection attempts to detect an injected signal and distinguish it from a signal originating from a human. This technique has been applied with varying degrees of effectiveness for other biometrics. The case of fingerprint readers in particular showcases that this is most likely an arms race between system designers considering more indications of liveness (e.g., the presence of skin oils or moisture for fingerprint readers) and attackers spoofing these indicators.

The second approach would be to keep the biometric data secret, thus attempting to prevent the attacker from obtaining any of the victim's ECG data. One could make the case that the Nymi Band's SDK should not allow for ECG data to be recorded, either through NEAs or otherwise. Disabling this functionality would make it at least significantly harder for the attacker to directly obtain data with the correct ECG morphology. As creating the mapping function requires training data from the target device, making this data hard to obtain would also raise the bar to carry out the attack. However, it might still be possible to determine the device-specific effects on the signal by analysing the hardware, rather than through empirical analysis of recorded signals. As we have demonstrated, the mapping function allows the attacker to obtain data from an arbitrary device and use it to attack the authentication system. While the probability of success somewhat depends on the device, we have demonstrated successful attacks using all of the devices we analyzed. Due to the large variety of devices recording ECG (e.g., medical ECG monitors, fitness devices, e-health) for purposes other than authentication, we do not consider keeping the ECG data secret a viable strategy.

In this work we have applied the mapping function in a way that allows us to attack an ECG-based authentication system with data collected on a different device. However, the same approach can be used to improve the interoperability of a biometric across different devices. For example, once a user is successfully enrolled for one device (e.g., a smartphone for the touchscreen input biometric), he could be authenticated on any other device without the need for re-enrolment, provided a mapping function between these devices has been trained beforehand.

IX. Conclusion

In this work we have presented a novel presentation attack against ECG biometrics and have demonstrated its effectiveness in attacking the Nymi Band. The Nymi Band is a commercially available wristband that is currently being trialled by MasterCard to authenticate the wearer for contactless payments. After collecting ECG data from 41 volunteers on different devices, we observed that the distributions of biometric features vary between these devices as well as different ECG measurement positions (Leads). This presents an attacker with difficulties when performing a presentation attack with data gathered on a different device. To address this challenge we describe a so-called mapping function, which transforms ECG signals measured on any device to change its morphology to match that of the Nymi Band. We use half of the subjects in our dataset to train the mapping function while the remainder is set aside to evaluate the actual attack. We use a number of signal injection methods, including a hardware signal generator and the playback of an ECG signal encoded as a .wav file. As such, the technological barriers to carry out the attack are extremely low in terms of both cost and hardware requirements. Our results show that injecting ECG data collected with the Nymi Band shows the highest success rate (81%). When using only the remaining signal sources we achieve a success rate of 43%, which increases to 62% after applying the mapping function.

Our results confirm the seriousness of the attack. Due to the universality of the mapping function and our signal injection devices, we expect the attack to show similar performance against other ECG-based authentication systems. We intend to further explore this once more systems become commercially available.

References

[1] F. Agrafioti, D. Hatzinakos, and J. Gao, *Heart biometrics: Theory, methods and applications*. INTECH Open Access Publisher, 2011.

[2] B. Barbot, M. Kwiatkowska, A. Mereacre, and N. Paoletti, "Estimation and verification of hybrid heart models for personalised medical and wearable devices," in *Computational Methods in Systems Biology*. Springer, 2015, pp. 3–7.

[3] J. Bonneau, "The science of guessing: analyzing an anonymized corpus of 70 million passwords," in *2012 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2012, pp. 538–552.

[4] G. D. Clifford, F. Azuaje, and P. McSharry, "Ecg statistics, noise, artifacts, and missing data," *Advanced Methods and Tools for ECG Data Analysis*, vol. 6, p. 18, 2006.

[5] G. W. Corder and D. I. Foreman, *Nonparametric statistics: A step-by-step approach*. John Wiley & Sons, 2014.

[6] S. Z. Fatemian, F. Agrafioti, and D. Hatzinakos, "Heartid: Cardiac biometric recognition," in *Biometrics: Theory Applications and Systems (BTAS), 2010 Fourth IEEE International Conference on*. IEEE, 2010, pp. 1–5.

[7] M. Frank, R. Biedert, E. Ma, I. Martinovic, and D. Song, "Touchalytics: On the applicability of touchscreen input as a behavioral biometric for continuous authentication," *IEEE transactions on information forensics and security*, vol. 8, no. 1, pp. 136–148, 2013.

[8] A. Fratini, M. Sansone, P. Bifulco, and M. Cesarelli, "Individual identification via electrocardiogram analysis," *Biomedical engineering online*, vol. 14, no. 1, p. 1, 2015.

[9] D. E. Goldberg, *Genetic algorithms*. Pearson Education India, 2006.

[10] F. E. Grubbs, "Sample criteria for testing outlying observations," *The Annals of Mathematical Statistics*, pp. 27–58, 1950.

[11] D. M. Hawkins, "The problem of overfitting," *Journal of chemical information and computer sciences*, vol. 44, no. 1, pp. 1–12, 2004.

[12] Z. Li, W. Han, and W. Xu, "A large-scale empirical analysis of chinese web passwords," in *Proc. 23rd USENIX Security Symposium, USENIX Security (August 2014)*, 2014.

[13] J. Ma, W. Yang, M. Luo, and N. Li, "A study of probabilistic password models," in *Security and Privacy (SP), 2014 IEEE Symposium on*. IEEE, 2014, pp. 689–704.

[14] J. Malmivuo and R. Plonsey, *Bioelectromagnetism: principles and applications of bioelectric and biomagnetic fields*. Oxford University Press, USA, 1995.

[15] M. L. Mazurek, S. Komanduri, T. Vidas, L. Bauer, N. Christin, L. F. Cranor, P. G. Kelley, R. Shay, and B. Ur, "Measuring password guessability for an entire university," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 173–186.

[16] P. E. McSharry, G. D. Clifford, L. Tarassenko, L. Smith *et al.*, "A dynamical model for generating synthetic electrocardiogram signals," *Biomedical Engineering, IEEE Transactions on*, vol. 50, no. 3, pp. 289–294, 2003.

[17] S. J. Orfanidis, *Introduction to signal processing*. Prentice-Hall, Inc., 1995.

[18] A. Pantelopoulos and N. G. Bourbakis, "A survey on wearable sensor-based systems for health monitoring and prognosis," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 40, no. 1, pp. 1–12, 2010.

[19] R. Sameni, M. B. Shamsollahi, C. Jutten, and G. D. Clifford, "A nonlinear Bayesian filtering framework for ECG denoising," *Biomedical Engineering, IEEE Transactions on*, vol. 54, no. 12, pp. 2172–2185, 2007.

[20] A. Serwadda and V. V. Phoha, "When kids' toys breach mobile phone security," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 599–610.

[21] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter, "Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1528–1540.

[22] E. M. Stein and R. Shakarchi, *Functional analysis: introduction to further topics in analysis*. Princeton University Press, 2011, vol. 4.

[23] E. Stobert and R. Biddle, "The password life cycle: user behaviour in managing passwords," in *Proc. SOUPS*, 2014.

[24] F. Sufi, I. Khalil, and J. Hu, "Ecg-based authentication," in *Handbook of Information and Communication Security*. Springer, 2010, pp. 309–331.

[25] C. M. Tey, P. Gupta, and D. Gao, "I can be you: Questioning the use of keystroke dynamics as biometrics," 2013.

[26] H. Wimberly and L. M. Liebrock, "Using fingerprint authentication to reduce system security: An empirical study," in *2011 IEEE Symposium on Security and Privacy*. IEEE, 2011, pp. 32–46.

[27] N. Zheng, K. Bai, H. Huang, and H. Wang, "You are how you touch: User verification on smartphones via tapping behaviors," in *2014 IEEE 22nd International Conference on Network Protocols*. IEEE, 2014, pp. 221–232.