On the Vulnerability of ECG Verification to Online Presentation Attacks

Nima Karimian
University of Connecticut
Electrical & Computer Engineering
nima@engr.uconn.edu

Damon L. Woodard and Domenic Forte
University of Florida
Electrical & Computer Engineering
dwoodard, dforte@ece.ufl.edu

Abstract

Electrocardiogram (ECG) has long been regarded as a biometric modality which is impractical to copy, clone, or spoof. However, it was recently shown that an ECG signal can be replayed from arbitrary waveform generators, computer sound cards, or off-the-shelf audio players. In this paper, we develop a novel presentation attack where a short template of the victim's ECG is captured by an attacker and used to map the attacker's ECG into the victim's, which can then be provided to the sensor using one of the above sources. Our approach involves exploiting ECG models, characterizing the differences between ECG signals, and developing mapping functions that transform any ECG into one that closely matches an authentic user's ECG. Our proposed approach, which can operate online or on-the-fly, is compared with a more ideal offline scenario where the attacker has more time and resources. In our experiments, the offline approach achieves average success rates of 97.43% and 94.17% for non-fiducial and fiducial based ECG authentication. In the online scenario, the performance is degraded by 5.65% *for non-fiducial based authentication, but is nearly unaffected for fiducial authentication.*

1. Introduction

As the Internet of things (IoT) becomes more popular in consumer, business, and military settings, one can expect the demand for biometric technologies to grow. IoT devices are supposed to outnumber the world's population this year [1] and their number should continue to dramatically increase for many years to come. Managing so many devices with passwords alone is ripe with challenges. In addition, the sensitive data gathered and stored by IoT could pose significant privacy concerns. Compared with conventional authentication techniques, such as digital passwords, personal identification numbers, and smartcards/tokens, biometrics provide a more robust method for identifying a person, i.e., based on their distinctive physical characteristics. Biometrics can also be considered more seamless and convenient, especially for continuous authentication [15]. That being said, it has already been demonstrated that many of the most popular biometric modalities (iris, face, fingerprint, and speech) can be spoofed and are, therefore, vulnerable to presentation attacks [27, 12, 13, 9, 3, 6, 28, 25, 24, 11]. Over the past decade, alternative modalities based on biological signals have been explored and their resistance to presentation attacks is often highlighted as a major attribute. Notable examples include electrocardiogram (ECG) [8], photoplethysmogram (PPG) [16], and electroencephalogram (EEG) [19], which possess high distinctiveness, are difficult to replicate, and provide intrinsic liveness detection. Among them, ECG has received the most attention and is beginning to gain larger acceptance from the biometrics community. For instance, ECG-based authentication systems, such as the Nymi wristband [2], are already coming to the market. ECG is a recording of the electric potential, generated by the electric activity of the heart, on the surface of the thorax that represents the extra cellular electric behavior of the cardiac muscle tissue. A typical, healthy ECG signal with different beats is shown in Figure 1. Generally speaking, ECG authentication systems can be categorized based on the feature extraction method (fiducial point vs. non-fiducial point) as well as the type of template matching used for classification. Fiducial point feature extraction relies on an accurate

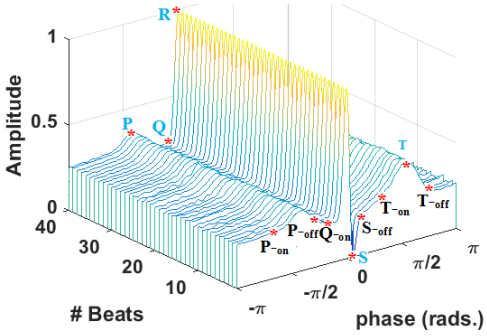

Figure 1. Waterfall plot of ECG beats collected from the same subject and localization of fiducial points.

978-1-5386-1124-1/17/$31.00 ©2017 IEEE

detection of ECG fiducial characteristic points such as P, Q, R, S, and T waves as shown in Figure 1, in order to obtain their relative amplitude, temporal intervals and morphological features. Non-fiducial point feature extraction analyzes an ECG in a holistic manner, typically by applying time or frequency analysis to obtain other statistical features. Despite the interest in ECG-based authentication, it's worth noting that ECG suffers from various noise sources such as motion, electromyography (EMG), and exercise, which can impact authentication accuracy [17]. In the literature, accuracy lies in the range of 94.3% to 100% [21].

Although ECG has long been considered as unclonable by many researchers [8, 21, 7, 18], that belief has been challenged recently. To the best of our knowledge, [10] is the first work to show how an ECG can be spoofed. Specifically, they tested a replay attack on the Nymi wristband by using three different types of devices to generate an ECG waveform: arbitrary waveform generators (AWGs), computer sound cards, and off-the-shelf audio players. The latter option is cheap, obtainable in a small form factor, and very effective. They achieve an 81% success rate when replaying the user's ECG via the above sources. The authors also consider a case where the ECG template in their possession is captured by a different biometric sensor other than Nymi's (e.g., at a physician's office). They develop a linear mapping function that transforms a signal recorded from one device, the source (e.g., physician), to a target device (e.g., Nymi). A 50% success rate is the best they were able to achieve when mapping from one source to another.

In this paper, we go beyond the above replay attack and aim for a full-fledged ECG presentation attack[1]. Instead of mapping from one source device to another, we consider mapping an attacker's ECG to the authentic user's ECG in order to falsely authenticate the attacker. Our approach exploits McSharry et al.'s [20] non-linear dynamical ECG model to accomplish this. Then, we generate a linear mapping between the models parameters based on the difference between fiducial features extracted from the source and target ECGs. There are three other important differences between this paper and [10]. First, our method requires only a single ECG beat (approximately one second) as a template rather than a long sequence of ECG signals (samples) to compute our mapping functions. Second, the mapping in [10] was calculated in an offline manner where time and hardware are unlimited. In this paper, we also consider an online scenario where the mapping needs to be computed on-the-fly with limited resources. The offline scenario is only used as a basis for comparison. Third, we consider different ECG feature extraction methods (fiducial and non-fiducial) and classification methods when evaluating the success rate of the proposed attack. Results show that the proposed presentation attack is successful more than 90% of the time on average in the worst case scenario (i.e., when only one beat of the victims ECG is available to execute the attack).

The remainder of the paper is organized as follows. In Section 2, we introduce the notation used throughout the paper and provide a high level overview of the proposed ECG presentation attacks. In Section 3, we discuss ECG signal and noise modeling as well as the online and offline proposed mapping approaches. Section 4 outlines our experimental setup and discusses simulation results. Finally, we conclude the paper in Section 5.

2. Overview of ECG Presentation Attack

The ECG presentation attack in the context of this paper can be described as follows. We assume that there exists an ECG-based biometric system with a legitimate user enrolled. The legitimate user's ECG signal is denoted by Y. We define ECG transformation as the process of learning a mapping function $F(\cdot)$. F takes as input an attacker's ECG signal (X) and the authentic user's signal (Y), and outputs a new ECG signal $\hat{Y}$ that is supposed to closely resemble Y. Note that we often refer to the legitimate user as the victim and Y as the victim's record. In order to capture different resource constraints of the attacker, we consider both online and offline attacks scenarios which are described in the subsections below.

The following notation is used for the remainder of the paper:

- A bold capital letter denotes a matrix (e.g., $\boldsymbol{A}$).
- A vector is represented by a lower case letter that is accented by a right arrow (e.g., $\vec{a}$).
- The ith element of a vector is denoted using a circular bracket notation (e.g., $a(i)$).
- We denote the attacker's ECG by X and the victim's ECG by Y. The attacker's ECG mapped to the victim's (an emulation of the victim's ECG) is denoted by $\hat{Y}$. A vector related to an ECG is denoted by the following modifier $(\vec{a}).)$. For example, the feature vector of the attacker is denoted by $(\vec{f})_X$.
- $\vec{c} = \frac{\vec{a}}{\vec{b}}$ and $\vec{c} = \vec{a} \times \vec{b}$ denote element-wise division and multiplication of vectors respectively (i.e., $c(i) = \frac{a(i)}{b(i)}$ and $d(i) = a(i)b(i)$).

2.1. Offline attack

An offline ECG presentation attack is illustrated in Figure 2(a). In this case, the mapping function F is determined using an expensive setup (e.g., PC or server) and there are no time constraints. The latter also implies that some elements of processing can be done manually (e.g., extraction

[1]Note that while the proposed approach is specific to ECG biometrics, the overall methodology might also be applicable to other electrophysiological signals such as EEG, PCG, and PPG as well.

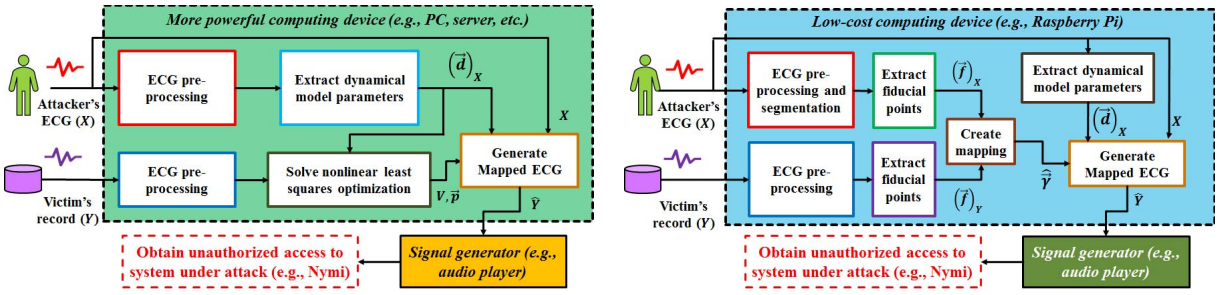

(a) (b)

Figure 2. Block diagrams of (a) offline and (b) online presentation attacks.

of fiducial features in the victim and attacker ECGs). The main steps are as follows. The attacker's ECG signal X and the victim's record Y are pre-processed, i.e., filtered to remove noise and segmented into beats. Then, modeling parameters of the attacker's ECG signal are extracted. In order to create a mapping function, an optimization problem is formulated and solved. In this paper, we minimize the squared Euclidean norm between the victim's record and attacker ECG signal (more details in Section 3.3).

The output of the optimization is a mapping function that transforms the dynamical model parameters of the attacker's ECG so that they resemble those of the victim's. A synthetic ECG signal is generated using McSharry et al.'s [20] non-linear dynamical ECG model and the transformed parameters. Synthetic noise is also mixed with the ECG in order to avoid simple presentation attack detection schemes by the biometric system. This signal $\hat{Y}$, an emulated version of the victim's ECG derived from the attacker's, can be stored in a low-cost device (e.g., audio player of a smart phone [10]) and later provided to an ECG sensor by the attacker in order to fool the biometric system. *Note that this offline attack is not necessarily a realistic or worthwhile attack. Instead, we use it to represent an idealized presentation attack to compare with our online version.*

2.2. Online attack

The online ECG presentation attack is illustrated in Figure 2(b). In the online case, the attacker's ECG is captured, pre-processed, and mapped per segment using a low-cost hardware platform (e.g., raspberry Pi). After pre-processing, fiducial point and temporal features are extracted from the victim ECG and each segment of the attacker's ECG. Instead of solving an optimization problem, the fiducial features are compared and a simpler linear mapping function is computed. Dynamical modeling parameters are extracted from the attacker's ECG and then mapped to the victim's parameters using this function. This process is repeated for each segment of the attacker's ECG as it is measured. A synthetic ECG signal is generated similar to the offline case and played to the biometric sensor via a low-cost audio player.

Compared to the offline case, the online approach is simpler because it does not require all segments of the ECG and only computes a linear mapping function. In addition, we shall only use the fiducial features that are computationally easy to extract.

3. ECG Modeling Preliminaries and Mapping Function Generation

3.1. ECG Dynamic Model

We introduce an analytical model that considers instantaneous heart rate in order to align multiple ECG beats. The technique transforms the signal from the time domain to angular domain, where each beat starts at angle $\theta = -\pi$ and ends at $\theta = \pi$. Figure 1 shows an example of ECG beats that have been aligned in the transformed, angular domain.

We also adopt the non-linear dynamical model proposed by McSharry et al. [20] to extract parameters from an ECG and generate synthetic ECGs for the aforementioned presentation attacks. McSharry et al.'s model uses three ordinary differential equations. It consists of a circular limit cycle of unit radius in the (x, y) plane around which the trajectory is pushed up and down as it approaches the P, Q, R, S and T points in the ECG:

$$\begin{cases} \frac{dx}{dt} = \beta x - \omega y \\ \frac{dy}{dt} = \beta y + \omega x \\ \frac{dz}{dt} = -\sum_{i \in P,Q,R,S,T} a_i \Delta\theta_i exp[-\frac{\Delta\theta_i^2}{2b_i^2}] - (z - z_0) \end{cases} \quad (1)$$

where $\beta = 1 - \sqrt{x^2 + y^2}$, $\Delta\theta_i = (\theta - \theta_i) mod(2\pi)$, $\theta = tan^{-1}(\frac{y}{x})$, the angular position of the elements of x,y range over $[-\pi, \pi]$, and ω is the angular velocity of the trajectory as it moves around the limit cycle. z_0 is the contribution from baseline wander and is assumed to be a relatively low frequency signal component coupled with the respiratory sinus frequency (RSA). The z axis represents the dynamics of the cardiac signal for the set of different fiducial points where θ_i is the location of the fiducial (PQRST) points, ω

represents heart rate, a_i and b_i are amplitude and variance of fiducial points for model parameters respectively ($i \in \{P, Q, R, S, T\}$). The above dynamic state equations can also be transformed into polar coordinates as follows [26]

$$\begin{cases} \frac{dr}{dt} = r(1-r) \\ \frac{d\theta}{dt} = \omega \\ \frac{dx}{dt} = -\sum_{i \in P,Q,R,S,T} a_i \Delta\theta_i exp[-\frac{\Delta\theta_i^2}{2b_i^2}] - (z - z_0) \end{cases} \quad (2)$$

where r and θ are the radial and angular state variables in polar coordinates. The transformation of Eq.(1) to Eq.(2) makes the second and third equations (2) to be independent from r. Since we align an ECG signal by individual beats, we can eliminate the baseline component $z - z_0$. This leaves the dynamics of Z as only a simple derivative of a sum of Gaussians, and it is possible to get an analytical solution of Z

$$Z(\theta) = -\sum_{i \in P,Q,R,S,T} \alpha_i exp(-\frac{(\theta - \theta_i)^2}{2b_i^2}) \quad (3)$$

where $\alpha_i = \frac{a_i b_i^2}{\omega}$ are the peak amplitude of the Gaussian functions used for modeling each of the ECG components. The analytical solution reduces $z(t)$ to $Z(\theta)$, as a given ECG beat with known angular position $\theta_k = \omega t_k$. $z(\theta)$ is basically a sum of Gaussian functions with means of each Gaussian at θ_i (fiducial point locations of the $PQRST$ complex).

3.2. ECG Noise Model

The three main types of noise sources in raw ECG signals are (1) motion artifacts (MA) which occur due to poor contact to the sensor; (2) baseline wander (BW) caused by body movement; and (3) electromyography (EMG) due to electrical activity of muscles, which is often non-stationary in time. In the context of this paper, noise impacts ECG authentication as well as our ability to extract dynamical parameters, fiducial points, and determine optimal mapping functions. Hence, it is important to have an ECG model with the flexibility to add/remove noise before and after mapping.

We use a time-varying auto-regressive (AR) parametric model to learn the noise parameters from the attacker's ECG signal. For the discrete time series of noise $y(n)$, a time-varying AR model of order p can be written as follows

$$y(n) = -\sum_{i=1}^{p} a_n(i) y(n-i) + e(n) \quad (4)$$

where $e(n)$ is the observation error and coefficients $a_n(i)$ ($i = 1, \ldots, p$) are the p time-varying AR parameters at the time instance of n. By defining $\vec{\eta}_n = [a_n(1), a_n(2), \ldots, a_n(p)]$ as a state vector, and $\vec{h}_n = [y(n-1), y(n-2), \ldots, y(n-p)]$ as the observation model, we can formulate the problem of AR parameter estimation in the Kalman Smoother (KS) form

$$y(n) = \vec{h}_n \vec{\eta}_n^T + e(n) \quad (5)$$

The progress of the state (i.e., the AR parameters) $\vec{\eta}_n$ when no prior information is available is typically described by a random walk model

$$\vec{\eta}_{n+1} = \vec{\eta}_n + \vec{\omega}_n \quad (6)$$

where $\vec{\omega}_n$ is the state noise term. Equations (5) and (6) form the state-space signal model for the time-varying AR process $y(n)$ and the evaluation of the AR parameters can now be estimated by using the Kalman Smoother algorithm [14].

Note that when determining the mapping function, the victim's ECG record and attacker's ECG are both filtered to remove noise. Once the mapping function is ready, a noise-free synthetic ECG is generated using the dynamical model and mapped parameters. Since ECG signals are inherently noisy, a simple anti-spoofing technique would likely be able to detect a presentation attack due to lack of noise in the synthetic ECG. Hence, the noise parameters acquired by KF smoothing are used to regenerate noise and mix it with the synthetic ECG.

3.3. Mapping Function Creation and Application

In this section, we illustrate the process of mapping ECG signals from the attacker to the victim. Since the victim ECG signal corresponds to the one enrolled in the biometric system, the attacker's mapped signal must be transformed in order to execute a presentation attack. To this end, we consider both online and offline attack scenarios for creating and applying the mapping function to mimic the victim's ECG signal.

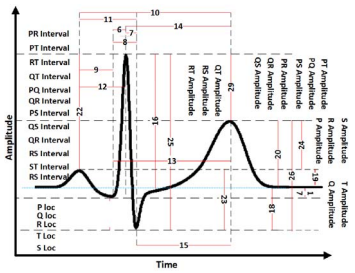

(a)

Extracted Fiducial Features

1-Q	14-RT	27-QS_amp
2-R	15-TS	28-RQ_amp
3-S	16-R_amp	29-RT_amp
4-P	17-Q_amp	30-P_on
5-T	18-S_amp	31-P_off
6-RQ	19-P_amp	32-Q_on
7-SR	20-T_amp	33-S_off
8-SQ	21-PQ_amp	34-T_on
9-QP	22-RP_amp	35-T_off
10-TP	23-PS_amp	
11-SP	24-PT_amp	
12-RP	25-RS_amp	
13-TQ	26-TS_amp	

(b)

Figure 3. (a) Electrocardiogram (ECG) PQRST complex and fiducial characteristic points, (b) 35 fiducial feature extracted for each ECG beat.

(1) Online mapping function: Our online mapping function is based on an observation that there is a linear relationship between fiducial features of an ECG signal and the dynamical model parameters from Eq.(3). Figure 3 shows a

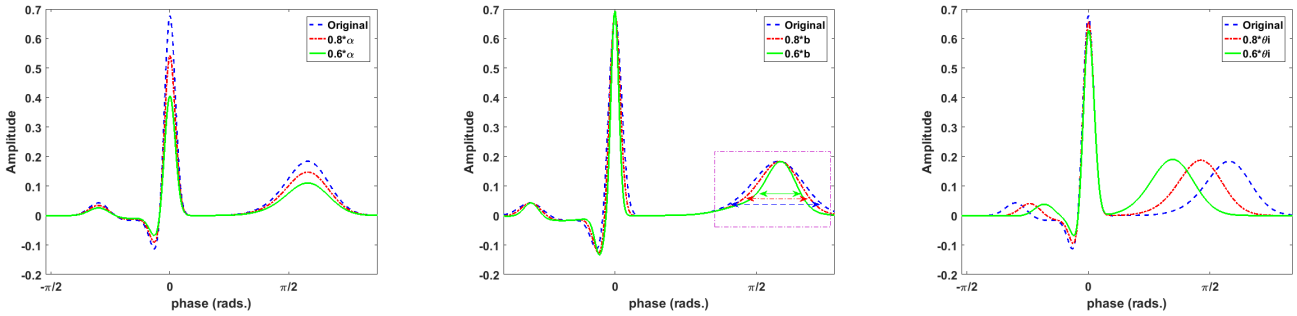

(a) (b) (c)

Figure 4. Impact of ECG signal by changing dynamical model parameters: (a) decreasing α parameters, (b) decreasing b parameters, and (c) decreasing θ parameters.

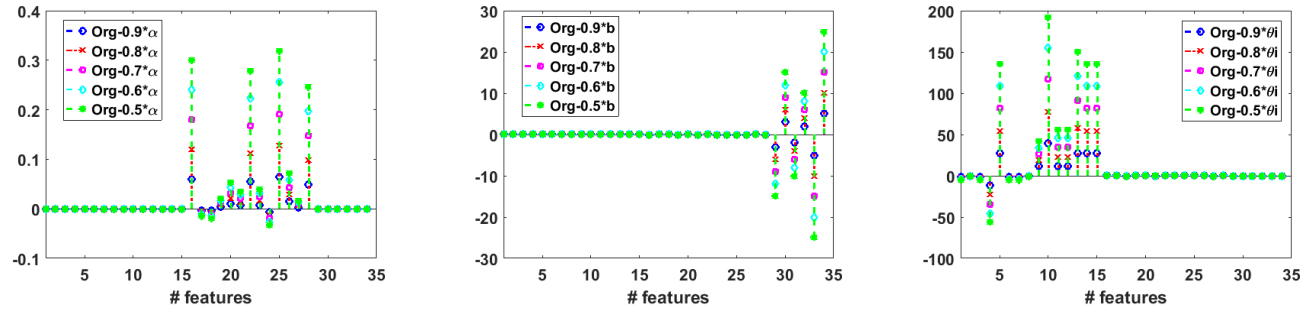

(a) (b) (c)

Figure 5. Impact of fiducial ECG features by changing dynamical model parameters: (a) decreasing α parameters, (b) decreasing b parameters, and (c) decreasing θ parameters.

set of 35 fiducial features that are often extracted when analyzing ECGs. For reasons that will be clear later, we divide them into three classes: 1) fiducial points excluding onsets and offsets (#1 to #15), 2) fiducial point amplitudes (#16 to #29), and 3) onsets and offsets (#30 to #35).

To uncover the above observation, we vary the dynamical model parameters by type ($\theta_i, \alpha_i, b_i \forall i$) and analyze the impact on each fiducial feature. The results are shown in Figures 4 and 5. In both figures, each dynamical model parameter is scaled by a factor ($0.9-0.5$). Figures 4(a-c) show how the ECG changes when scaling α, b, and θ parameters respectively. Figure 5(a-c) shows how scaling these same parameters changes the fiducial points. As shown in Figure 5 (a), the amplitudes (features 16 to 29) possess a linear relationship with the α parameters. In contrast, the onset and offset of fiducial point ECG features (features 30 to 35) are translated linearly by scaling factors associated with b parameters (see Figure 5(b)). Finally, Figure 5 (c) shows a similar relationship between θ parameters and the PQRST complex fiducial points (features 1 to 15).

Based on the above observation, a simple mapping function can be constructed from a single heartbeat (segment) of X using the following methodology. Assume fiducial features are extracted from victim's record Y (assumed to be one segment for simplicity) and a segment from X. We represent a feature vector as $\vec{f} = [\vec{f_\theta}, \vec{f_\alpha}, \vec{f_b}]$ where $f_\theta(i)$, $f_\alpha(i)$, and $f_b(i)$ correspond to features 1-15, 16-29, and 30-35 respectively. Also, we denote the victim and attacker feature vectors by $(\vec{f})_Y$ and $(\vec{f})_X$ respectively. Similarly, there also exist vectors of dynamical model parameters represented by $\vec{d} = [\vec{d_\theta}, \vec{d_\alpha}, \vec{d_b}]$ where $\vec{d_u} = [u_p, u_Q, u_R, u_S, u_T]^T, u \in \{\theta, \alpha, b\}$. Note that based on the above discussion, it is convenient to divide both $\vec{f}$ and $\vec{d}$ in this manner due to the relationship between θ parameters and features 1-15, α parameters and features 16-20, and so forth.

We define the scaling vectors $\vec{\tau} = \frac{(\vec{f})_Y}{(\vec{f})_X}$ and $\vec{\gamma} = \frac{(\vec{d})_Y}{(\vec{d})_X}$. As discussed above, there is a linear relationship between these scaling factors

$$\vec{\gamma} = \boldsymbol{S}\vec{\tau} + \vec{o} \tag{7}$$

$$\boldsymbol{S} = \begin{pmatrix} \boldsymbol{A}_{(5\times15)} & \boldsymbol{0}_{(5\times14)} & \boldsymbol{0}_{(5\times6)} \\ \boldsymbol{0}_{(5\times15)} & \boldsymbol{B}_{(5\times14)} & \boldsymbol{0}_{(5\times6)} \\ \boldsymbol{0}_{(5\times15)} & \boldsymbol{0}_{(5\times14)} & \boldsymbol{C}_{(5\times6)} \end{pmatrix} \tag{8}$$

where matrices $\boldsymbol{A}$, $\boldsymbol{B}$, and $\boldsymbol{C}$ relate fiducial scaling factors to scaling factors for dynamical parameters θ, α, and b respectively. Similarly, $\vec{o}$ is an offset.

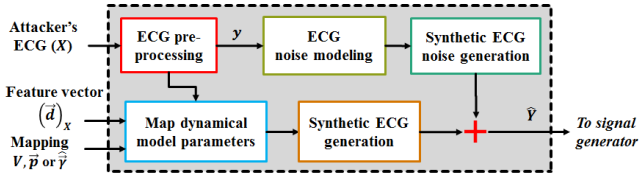

Figure 6. Application of mapping function for presentation attack.

The above implies that, given $(\vec{f})_Y$, $(\vec{f})_X$, and $(\vec{d})_X$, one can compute $(\vec{d})_{\hat{Y}}$ as follows

$$\begin{aligned}(\vec{d})_Y &= \vec{\gamma} \times (\vec{d})_X \\ (\vec{d})_{\hat{Y}} &\approx (\boldsymbol{S}\vec{\tau} + \vec{o}) \times (\vec{d})_X \\ (\vec{d})_{\hat{Y}} &\approx \left(\boldsymbol{S}\frac{(\vec{f})_Y}{(\vec{f})_X} + \vec{o}\right) \times (\vec{d})_X\end{aligned} \tag{9}$$

The above mapping method relies on a linear transformation between sets of fiducial ECG features from one beat, and can be easily implemented on very low-cost hardware (e.g., raspberry pi). In practice, however, it is difficult to extract features 30 to 35 (corresponding to offsets and onsets). *Thus, in our results section, we base our mapping on features 1 to 29 in order to compute* $\vec{d}_{\hat{Y}}$. Note that the above approach operates on a single segment (heartbeat) of the attacker. Therefore, we must apply it to each segment of X separately in our later experiments.

(1) Offline mapping function: The online approach is limited by the fact that $(\vec{d})_{\hat{Y}}$ is computed using a linear transform and does not include b fiducial features. To deal with this issue, we formulate a non-linear optimization problem that finds a better mapping without the need to extract fiducial features. While it is more flexible and accurate, we may only apply it in offline scenarios because it requires more time and segments of X to compute.

The optimization problem is given as follows:

$$[V^*, \vec{p}^*] = \arg\min \sum_{\theta=-\pi}^{\pi} ||Y(\theta) - Z(\theta; (\vec{d})_{\hat{Y}}))||^2, \tag{10}$$

where $||\cdot||$ denotes the Euclidean norm and $*$ denotes the parameters of the optimal solution. In our results section, we use a linear mapping function $(\vec{d})_{\hat{Y}} = V(\vec{d})_X + \vec{p}$, but nonlinear functions can also be used in this formulation and shall be investigated in future work. The Levenberg-Marquardt algorithm was used to solve the non-linear least-squares optimization problem.

Once a mapping function (either online or offline) is obtained, we can apply it to the attacker's ECG signal X to generate spoofed ECG signal $\hat{Y}$. Application of the mapping function to the signal is shown in Figure 6. First, the attacker's signal is pre-processed by employing a 4^{th} order of Butterworth band pass filter with cutoff frequency 1Hz-40Hz. After that, the mapping function is applied to the filtered ECG signal in order to obtain new dynamic parameters. Based on the new parameters, a synthetic ECG signal is generated using McSharry et al.'s model (see Section 3.1). Meanwhile, the noise in X is modeled as described in Section 3.2 and then used to generate symmetric noise which is mixed with the synthetic ECG. The final output is a synthetic ECG signal $\hat{Y}$ with noise that is meant to emulate the victim's signal.

4. Experimental Results

In this section, we present the results obtained from our proposed attack with two popular feature extraction techniques.

4.1. Experimental Setup

ECG Data: The ECG recordings of 52 subjects from the Physikalisch-Technische Bundesanstalt (PTB) database [22] are selected and a template database is prepared. Each signal is digitized at 1000 samples per second, with 16 bit resolution and an average of two minutes using a single lead ECG. Note that the reason we have decided to only consider PTB database is because other ECG databases have different characteristics such as input resistance, input voltage, resolution, bandwidth and sampling rate. However, our proposed approaches can be applied to other ECG databases as well.

Feature Extraction: Two popular feature extraction techniques are applied in this paper

(1) Fiducial feature extraction: A subset of 29 features that represent the majority of fiducial features are extracted from every beat of each individual's ECG signal. As shown in figure 3, features encompass 21 fiducial points and 14 temporal features. To extract these features, first the R peak and then the P, Q, S, T peaks and valleys are detected using a local maximum/minimum searching algorithm within a defined physical region. However, note that feature numbers 30 to 35 are not considered in mapping or classification because they are difficult to extract.

(2) Normalized Autocorrelation (AC): The motivation behind this non-fiducial approach is the use of normalized autocorrelation (AC) method on non-overlapping windows of the filtered individual ECG signal without the use of fiducial point detection [23]. In order to extract the feature vector representing the ECG's signature, a windowed ECG signal and estimation of the normalized AC over a window of m are taken into account. In fact, autocorrelation gives an automatic shift invariant feature set that represents repetitive characteristics over multiple heartbeat cycles. The autocorrelation coefficients can be written as

$$\hat{R}_{xx}(m) = \frac{1}{\hat{R}_{xx}(0)} \sum_{i}^{N-|m|-1} s(i)s(i+m) \tag{11}$$

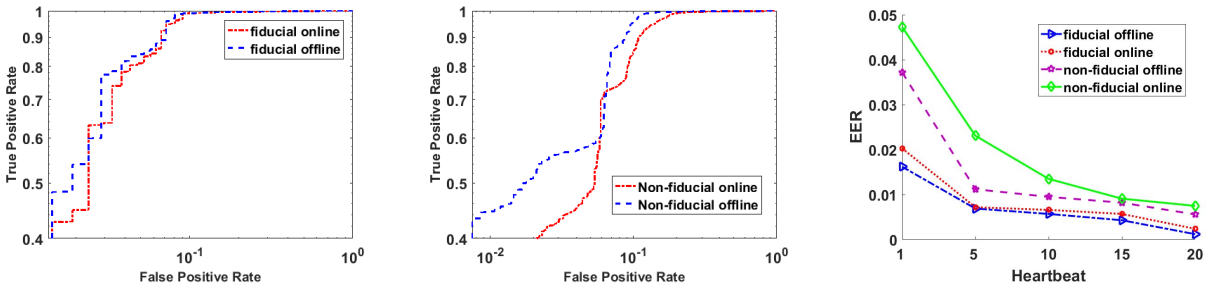

(a) (b) (c)

Figure 7. (a) ROC curves (log-log scale) for fiducial feature extraction under both online and offline attack; (b) ROC curves for nonfiducial feature extraction under both online and offline attacks; (c) EER based on training the mapping function for different number of victim heartbeats in all attack scenarios.

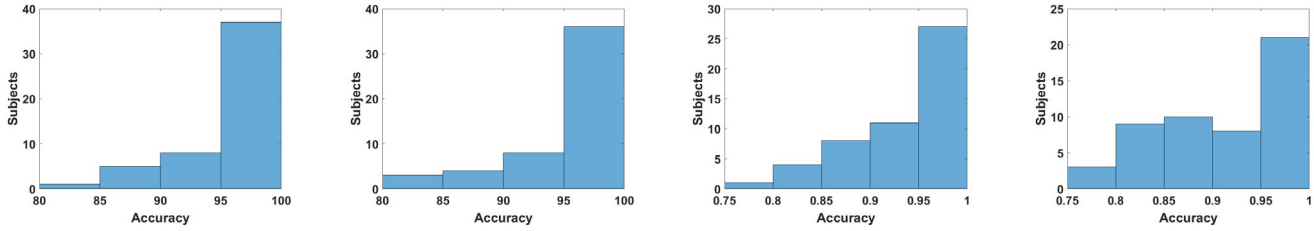

(a) (b) (c) (d)

Figure 8. Distribution of accuracy across subjects for fiducial feature extraction in (a) offline and (b) online modes; distribution of accuracy for non-fiducial feature extraction in (c) offline and (d) online modes.

where $s(i)$ is the windowed ECG signal at time i, $s(i+m)$] is the time shifted version of the windowed ECG with a time lag of m which is greater than the mean QRS duration

Classification: We apply a one-class support vector machine (SVM) technique for classification. The basic training principal of SVM is to find the optimal hyper-plane that separates the classes with a maximum margin [4]. In order to train the SVM, we use 40 different test samples (heartbeats) for any victim's ECG. In the results below, we take every subject as an attacker and consider the rest as victims. In other words, we apply the attack $52 \times 51 = 2,652$ times. Unless otherwise specified, one can assume the presented results represent an average. The libSVM library [5] is used for our experiments.

Experiment Parameters: We consider three cases. In case I, the victim's record is only one beat long and one emulated beat is generated for authentication. In case II, we increase the number of samples (beats) contained in the victim's record while still using only one beat to authenticate. In case III, we assume the victim's record is one beat long, but increase the number of samples used to authenticate the attacker. Three error rates are used to evaluate the attack performance: false positive/accept rate (FPR), true positive/accept rate (TPR) and equal error rate (EER). FPR is the percentage of attackers who were denied access to the system whereas TRP is the percentage of attackers who have successfully gained access to the system. The two error rates FPR and TPR can be traded-off with each other in order to find the optimal and desired EER. EER is the location on the receiver operator characteristic (ROC) curve where the FPR and TPR are equal. We also calculate the accuracy for each subject as the number of successful attempts by the attacker divided by the total number of attempts.

4.2. Case I

ROC curves are shown in Figure 7 (a-b) for the case where only one beat of victim's ECG signal is used to create the mapping function. As shown in this figure, the average accuracy of fiducial feature extraction for both online and offline attacks is 96.69% versus 97.43% while performing five-fold cross validation. In contrast, non-fiducial feature extraction obtains 91.78% and 94.17% rate of success for online and offline attacks. As expected, the offline attacks perform better than online attacks in terms of accuracy because the online mapping does not involve the entire ECG (ignores b parameters). However, it should be noted that the online attack's performance is quite comparable to the more advanced offline attack. In addition, the proposed approach has better success for fiducial feature extraction compared with non-fiducial. The characteristics of fiducial features (e.g. temporal and amplitude) generally make them more distinguishable compared to non-fiducial; thus higher accuracy is expected. This is true for both offline and online attacks.

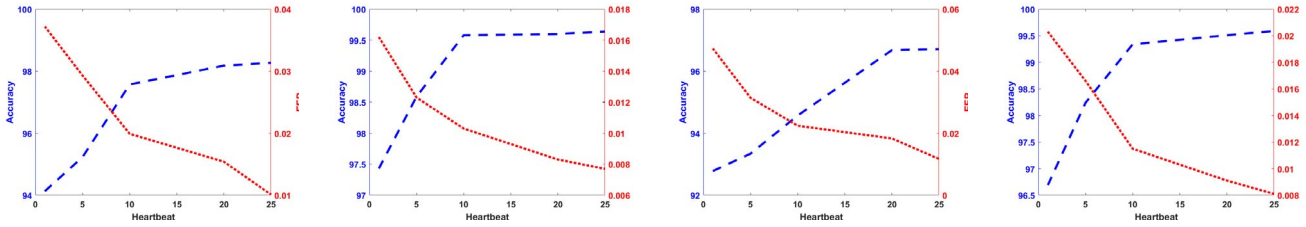

(a) (b) (c) (d)

Figure 9. Accuracy (red) and EER (blue) vs. numbers of heartbeats used during authentication. (a) & (b) show a comparison for offline attacks using non-fiducial and fiducial feature extraction, respectively; (c) & (d) show a comparison for online attacks using non-fiducial and fiducial feature extraction respectively.

In order to analyze the performance of authentication, we also study the accuracy of different subjects. The distribution of accuracy percentage across subjects for different attacks is shown in Figures 8(a-d). The average accuracy for fiducial feature extraction are 97.43% and 96.69% for offline and online cases respectively; while for non-fiducial accuracies are 94.17% and 91.78%. We find that some subjects are more difficult to execute presentation attacks on than others. This can be explained by two reasons. First, when only one beat of the victim's signal is used to generate the mapping function, the selected beat we've chosen might contain more noise compared to other subjects. Second, the noise mixed with the signal could also impact the authentication accuracy. Although we have added noise to the synthetic ECG (in order to avoid trivial detection of the presentation attack), the noise is modeled from the attacker's ECG signal instead of the victim's since we have no way of emulating the intra-beat variation/noise of victim's ECG signal (i.e., heart rate variability). By investigating the subjects with higher accuracy and lower accuracy, this hypothesis was confirmed; Attacker ECGs with similar in heart rate and heart rate variability to the victim's have a higher accuracy. In addition, we also note that subjects with lower accuracy for fiducial methods are correlated with the subjects with lower accuracy for non-fiducial methods (not explicitly shown for brevity).

4.3. Case II

To overcome the above heart rate variability issue, we also investigate the impact of mapping when multiple ECG beats of the victim are available for mapping. In Figure 7(c), the number of heartbeats (samples) is varied from 1 to 20. It can be observed that the EER decreases as the number of samples increases, which agrees with what we expect based on the above discussion. For non-fiducial feature extraction, the worst EERs for offline and online attacks (0.0372 & 0.0473) occur when the victim's record has only 1 beat. This improves by a factor of approximately 6 for 20 beats. Furthermore, the EERs for fiducial feature extraction improve by factors of 13 and 8 respectively when increasing from 1 beat to 20.

4.4. Case III

Figures 9 (a-d) show the accuracy and EER versus number of samples used to authenticate the attacker. In general, the impact of noise on accuracy often lessens with more samples for most applications. The figure shows that this trend holds even in the case of our emulated attacker ECGs. As the length of the ECG used to authenticate the attacker increases, the EER falls and the accuracy increases before eventually saturating. We were be able to achieve the accuracy of 99.64% and 99.51% for fiducial feature extraction based on offline and online attacks receptively. 98.27% and 96.71% accuracy for non-fiducial feature are obtained based on offline and online attacks.

5. Conclusion

In this work, we have proposed a presentation attack on ECG-based biometric systems. To the best of our knowledge, this is the first work to explore the vulnerability of ECG biometric by applying a systematic mapping function that transforms any attacker's ECG signal to a victim's. We evaluate the presentation attack on both offline and online attacks for two popular feature extraction methods (fiducial and non-fiducial). In our experiments, the proposed online approach achieves success rates over 90% with limited training data. When more training samples of the victim ECG are available to the attacker, the success rate rises to over 96%. The performance of the resource constrained online approach is even comparable to the online approach. In future work, we plan on implementing the proposed online approach in real hardware and evaluating different nonlinear mapping functions for offline scenarios. In addition, we shall also consider cases where the victim's ECG has been recorded by a different source device.

6. Acknowlegements

This project was supported in part by US Army Research Office grant under award number grant W911NF-16-1-0321.

References

[1] Gartner says 8.4 billion connected. http://www.gartner.com/newsroom/id/3598917. 1

[2] Nymi corporate website, Apr 2017. http://www.nymi.com. 1

[3] A. Anjos and S. Marcel. Counter-measures to photo attacks in face recognition: a public database and a baseline. In *Biometrics (IJCB), 2011 international joint conference on*, pages 1–7. IEEE, 2011. 1

[4] B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152. ACM, 1992. 7

[5] C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011. 7

[6] C. Chen, A. Dantcheva, T. Swearingen, and A. Ross. Spoofing faces using makeup: An investigative study. In *2017 IEEE International Conference on Identity, Security and Behavior Analysis (ISBA)*, pages 1–8, Feb 2017. 1

[7] S. Y. Chun. Single pulse ecg-based small scale user authentication using guided filtering. In *Biometrics (ICB), 2016 International Conference on*, pages 1–7. IEEE, 2016. 2

[8] H. P. Da Silva, A. Fred, A. Lourenço, and A. K. Jain. Finger ecg signal for user authentication: Usability and performance. In *Biometrics: Theory, Applications and Systems (BTAS), 2013 IEEE Sixth International Conference on*, pages 1–8. IEEE, 2013. 1, 2

[9] T. de Freitas Pereira, A. Anjos, J. M. De Martino, and S. Marcel. Can face anti-spoofing countermeasures work in a real world scenario? In *Biometrics (ICB), 2013 International Conference on*, pages 1–8. IEEE, 2013. 1

[10] S. Eberz, A. Patané, N. Paoletti, M. Kwiatkowska, M. Roeschlin, and I. Martinovic. Broken hearted: How to attack ecg biometrics. In *The Network and Distributed System Security Symposium (NDSS)*, February 2017. 2, 3

[11] S. K. Ergünay, E. Khoury, A. Lazaridis, and S. Marcel. On the vulnerability of speaker verification to realistic voice spoofing. In *Biometrics Theory, Applications and Systems (BTAS), 2015 IEEE 7th International Conference on*, pages 1–6. IEEE, 2015. 1

[12] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia. From the iriscode to the iris: A new vulnerability of iris recognition systems. *Black Hat Briefings USA*, 2012. 1

[13] P. Gupta, S. Behera, M. Vatsa, and R. Singh. On iris spoofing using print attack. In *Pattern Recognition (ICPR), 2014 22nd International Conference on*, pages 1681–1686. IEEE, 2014. 1

[14] S. S. Haykin et al. *Kalman filtering and neural networks*. Wiley Online Library, 2001. 4

[15] N. Karimian, Z. Guo, M. Tehranipoor, and D. Forte. Highly reliable key generation from electrocardiogram (ecg). *IEEE Transactions on Biomedical Engineering*, 64(6):1400–1411, June 2017. 1

[16] N. Karimian, Z. Guo, M. Tehranipoor, and D. Forte. Human recognition from photoplethysmography (ppg) based on non-fiducial features. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4636–4640, March 2017. 1

[17] N. Karimian, F. Tehranipoor, Z. Guo, M. Tehranipoor, and D. Forte. Noise assessment framework for optimizing ecg key generation. In *2017 IEEE International Symposium on Technologies for Homeland Security (HST)*, pages 1–6, April 2017. 2

[18] M. Komeili, W. Louis, N. Armanfard, and D. Hatzinakos. Feature selection for nonstationary data: Application to human recognition using medical biometrics. *IEEE Transactions on Cybernetics*, PP(99):1–14, 2017. 2

[19] E. Maiorana, D. L. Rocca, and P. Campisi. On the permanence of eeg signals for biometric recognition. *IEEE Transactions on Information Forensics and Security*, 11(1):163–175, Jan 2016. 1

[20] P. E. McSharry, G. D. Clifford, L. Tarassenko, and L. A. Smith. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Transactions on Biomedical Engineering*, 50(3):289–294, 2003. 2, 3

[21] I. Odinaka, P.-H. Lai, A. D. Kaplan, J. A. O'Sullivan, E. J. Sirevaag, and J. W. Rohrbaugh. Ecg biometric recognition: A comparative analysis. *IEEE Transactions on Information Forensics and Security (T-IFS)*, 7(6):1812–1824, 2012. 2

[22] M. Oeff, H. Koch, R. Bousseljot, and D. Kreiseler. The ptb diagnostic ecg database. *National Metrology Institute of Germany, http://www. physionet. org/physiobank/database/ptbdb*, 2012. 6

[23] P. S. Raj, S. Sonowal, and D. Hatzinakos. Non-negative sparse coding based scalable access control using fingertip ecg. In *Biometrics (IJCB), 2014 IEEE International Joint Conference on*, pages 1–6. IEEE, 2014. 6

[24] A. Rattani and A. Ross. Automatic adaptation of fingerprint liveness detector to new spoof materials. In *Biometrics (IJCB), 2014 IEEE International Joint Conference on*, pages 1–8. IEEE, 2014. 1

[25] A. Roy, N. Memon, and A. Ross. Masterprint: Exploring the vulnerability of partial fingerprint-based authentication systems. *IEEE Transactions on Information Forensics and Security*, 12(9):2013–2025, 2017. 1

[26] R. Sameni, M. B. Shamsollahi, C. Jutten, and G. D. Clifford. A nonlinear bayesian filtering framework for ecg denoising. *IEEE Transactions on Biomedical Engineering*, 54(12):2172–2185, 2007. 4

[27] S. Venugopalan and M. Savvides. How to generate spoofed irises from an iris code template. *IEEE Transactions on Information Forensics and Security*, 6(2):385–395, 2011. 1

[28] Q. Zhao, A. K. Jain, N. G. Paulter, and M. Taylor. Fingerprint image synthesis based on statistical feature models. In *Biometrics: Theory, Applications and Systems (BTAS), 2012 IEEE Fifth International Conference on*, pages 23–30. IEEE, 2012. 1