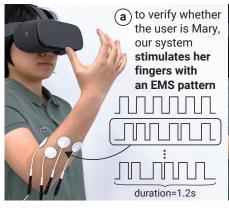
User Authentication via Electrical Muscle Stimulation

Yuxin Chen yxchen@cs.uchicago.edu University of Chicago Chicago, IL, USA

Pedro Lopes pedrolopes@cs.uchicago.edu University of Chicago Chicago, IL, USA Zhuolin Yang zhuoliny@cs.uchicago.edu University of Chicago Chicago, IL, USA

Ben Y. Zhao ravenben@cs.uchicago.edu University of Chicago Chicago, IL, USA Ruben Abbou rabbou@cs.uchicago.edu University of Chicago Chicago, IL, USA

Haitao Zheng htzheng@cs.uchicago.edu University of Chicago Chicago, IL, USA



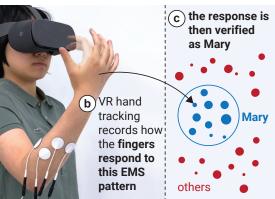




Figure 1: We propose a novel modality for authentication: electrical muscle stimulation (EMS). To explore it, we created an interactive system that (a) stimulates the user's forearm muscles with electrical impulses (i.e., using one of 68M possible EMS challenges); (b) measures the user's involuntary finger movements, which are unique because everybody's physiology is different; (c) verifies this response using an authentication model, and immediately eliminates this challenge, making our system secure against data breaches and replay attacks as it never reuses the same challenge. We demonstrate it here using the example of (d) authenticating a VR user without passwords or PINs.

ABSTRACT

We propose a novel modality for active biometric authentication: electrical muscle stimulation (EMS). To explore this, we engineered an interactive system, which we call ElectricAuth, that stimulates the user's forearm muscles with a sequence of electrical impulses (i.e., EMS challenge) and measures the user's involuntary finger movements (i.e., response to the challenge). ElectricAuth leverages EMS's *intersubject variability*, where the same electrical stimulation results in different movements in different users because everybody's physiology is unique (e.g., differences in bone and muscular structure, skin resistance and composition, etc.). As such, ElectricAuth allows users to login without memorizing passwords or PINs.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '21, May 8–13, 2021, Yokohama, Japan

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-8096-6/21/05...\$15.00 https://doi.org/10.1145/3411764.3445441

ElectricAuth's challenge-response structure makes it secure against data breaches and replay attacks, a major vulnerability facing today's biometrics such as facial recognition and fingerprints. Furthermore, ElectricAuth never reuses the same challenge twice in authentications – in just one second of stimulation it encodes one of 68M possible challenges. In our user studies, we found that ElectricAuth resists: (1) impersonation attacks (false acceptance rate: 0.17% at 5% false rejection rate); (2) replay attacks (false acceptance rate: 0.00% at 5% false rejection rate); and, (3) synthesis attacks (false acceptance rates: 0.2-2.5%). Our longitudinal study also shows that ElectricAuth produces consistent results over time and across different humidity and muscle conditions.

CCS CONCEPTS

• Human-centered computing → Human computer interaction (HCI); Haptic devices; • Security and privacy → Authentication; Biometrics.

KEYWORDS

electrical muscle stimulation, biometric authentication, wearable

ACM Reference Format:

Yuxin Chen, Zhuolin Yang, Ruben Abbou, Pedro Lopes, Ben Y. Zhao, and Haitao Zheng. 2021. User Authentication via Electrical Muscle Stimulation. In *CHI Conference on Human Factors in Computing Systems (CHI '21), May 8–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3411764.3445441

1 INTRODUCTION

Biometric authentication is a technique that identifies an individual by their unique biological characteristics, such as their iris [85], fingerprints [51], or even one's voice [8]. To identify their users, these interactive systems compare a previously stored biometric key to incoming, typically real-time, biometric data of the user wishing to authenticate. Compared to traditional password or PIN based systems, biometric authentication offers significantly better usability as it does not require users to memorize passwords or PINs. As such, biometric authentication is getting widely adopted, replacing passwords in many contexts [74].

However, the key feature of biometric authentication is typically also its key flaw: once the biometric data is compromised (e.g., stolen in database breaches or recorded by an external attacker), there is nothing the user can do to securely re-use their own data. For example, if someone steals a user's fingerprints, this user can never trust a fingerprint-based interactive system. Unfortunately, these threats are not theoretical and many biometric systems have been breached. For instance, the biggest known biometric data breach involved a database of 27.8M records, including fingerprints and faces [30].

To tackle this shortcoming, researchers turned to interactive systems that feature a *challenge-response* as a form of active biometric authentication. One example is Velody [40], which challenges a user by vibrating her palm and measuring the user's unique vibration-response. The advantage of these systems is that, if the stored challenge-response pairs are breached, the system can quickly recover by simply asking the user to submit responses to a new set of challenges. As such, researchers seek to find more modalities that afford challenge-response biometric authentication.

In this paper, we propose and explore a novel modality for active biometrics: electrical muscle stimulation (EMS). To understand and evaluate the potential of EMS as a biometrics system for interactive applications, we engineered a prototype that performs user authentication via EMS. Our system, which we call ElectricAuth, stimulates the wearer's forearm muscles with an EMS-based challenge, i.e., a 1.2s long sequence of electrical impulses on four of the user's muscles. Then, it measures the user's involuntary movements that result from this EMS challenge. In Figure 1, we illustrate our system with the example of authenticating a user in VR. Here, ElectricAuth uses the VR headset's hand tracking to observe the response of the user's muscles to the EMS-challenge as their individual finger muscles are actuated.

Electric Auth makes three key contributions in the design of EMS-based biometric authentication.

First, ElectricAuth authenticates users by leveraging what is typically seen as the biggest disadvantage of EMS: *intersubject variability*, i.e., the same electrical stimulation results in different movements in different users because everybody's physiology is different [11, 14, 17, 36, 53]. This unique response to EMS across users is well-known and well-documented in the early HCI works

that pioneered the use of EMS in interactive devices, for instance: "(...) stimulation level differed between users and was clearly dependent on the muscle and fat level and thickness of the arm" (from Kruijff et al. [39]) and, similarly, "(...) levels according to individual variations" (from PossessedHand [79]). In fact, researchers in the field of muscle-biomechanics and physiology demonstrated how this uniqueness arises from multiple factors, such as differences in muscle contractility [23], muscle elasticity [82], muscle viscosity [13], the limb's mass and shape [55], skin conductance [41], bioimpedance [12, 70] and even nerve conduction [1]. All these differences add up to create individual responses to the same stimulus, which our system uses as the key feature to authenticate a user.

Second, ElectricAuth generates a very large pool of challenges by exploring an underutilized property of EMS: muscles respond differently depending on their current state of contraction, which can be altered by varying the timing between two impulses. Using four muscles, six impulses and seven time gaps, ElectricAuth encodes one of 68M possible challenges in 1.2s. As such, ElectricAuth is robust against data breaches and replay attacks because it never reuses the same challenge twice in authentications – ElectricAuth rejects replay of recorded responses to any previously used challenges, and can quickly recover from leak/breach of either authentication model or stored challenge-response pairs by asking the user to register responses to a new set of challenges (like registering new one-time passwords).

Finally, we evaluated our prototype of ElectricAuth by means of four different evaluations, each shining light on a different facet of our research question: (1) in our user studies, we found that ElectricAuth offers accurate user verification and resists three common biometric attacks: impersonation, replay and synthesis attacks; (2) in our exploratory longitudinal study, we found that ElectricAuth's pre-trained authentication model performed stably over 21 days against various muscle conditions (fatigue, humidity, etc.) that were absent from the training data; (3) in our technical evaluation we showed that ElectricAuth, after receiving a response, can verify the user in 3ms on laptop's CPU and 35ms on a small embedded device; we also confirmed the use of depth camera as an alternative motion tracking modality (since our prototype uses IMUs); and, (4) we generated synthetic impersonator responses to test ElectricAuth's robustness against impersonation attacks at scales larger than our user studies.

2 RELATED WORK

The work presented in this paper builds on the fields of wearables, electrical muscle stimulation, and biometrics.

2.1 Electrical muscle stimulation

Electrical muscle stimulation (EMS) is a technique from medical rehabilitation [76] that induces involuntary movements by delivering electrical impulses to the user's muscles. This is typically achieved by non-invasive methods such as attaching pairs of electrodes to the user's skin (e.g., on top of the muscles that control finger movement, located in the forearm). Electrode pairs are typically driven using safe and medical compliant muscle stimulators [37].

The range of motion of an induced muscle contraction depends on several key factors. Even in the very first interactive use of EMS in HCI, by Kruijff et al. [39] in 2006, the potential causes of EMS' intersubject-variability were discussed: "(..) stimulation level differed between users and was clearly dependent on the muscle and fat level and thickness of the arm (...)". Similarly, in Possessed-Hand [79], Tamaki et al. also found "(...) stimulation levels according to individual variations". In fact, researchers in the fields of muscle-biomechanics and physiology have been investigating precisely which factors drive a muscle's unique response to electrical impulses, including: the location of the electrodes [68, 79]; the electrical waveform characteristics, such as frequency and amplitude of the impulses [39, 68, 79]; the target muscle's contractility [23], i.e., the ability of muscle fibers to shorten; muscle elasticity [26, 82], i.e., the ability of the elastic tissue present in the muscle fibers to return to its original length when a tensile force is removed; muscle viscosity [13], i.e., the internal bio-lubrication of the muscle inhibits the muscle from reacting too quickly to protect against stretch injuries; the limb's mass and shape [10, 11, 14, 17, 55]; skin conductance affects non-constant current EMS devices [39, 41], bioimpedance [12, 70]; and, even nerve conduction [1, 75], i.e., the speed of nerve signal transmission. However, it is not possible to precisely determine how much each factor weighs in the final variability, as these are tied together in complex non-linear ways, and this is still an open research question in muscle physiology. More importantly, all the aforementioned factors are relevant to our proposed technique since these vary-across users. Typically, a combination of these explains the intersubject variability seen in EMS-based interactive systems, which is why researchers report long periods of calibration [44, 47, 77, 79] and even specifically mention differences across users [39, 79].

Recently, researchers started to engineer interactive devices based on EMS. These tend to fall into two broad categories: (1) haptic devices that increase immersion/realism of virtual environments, and (2) interactive devices that facilitate information access via proprioception. As far as interactive devices that increase immersion, EMS has been used to render forces in mobile devices [43], virtual reality [44, 47] or augmented reality [20, 48]. As a means of general information access, EMS has been especially used for haptic training (e.g., learning a musical instrument [79], operating a tool the user is not familiar with [46]) or eyes-free communication (e.g., communicating walking directions via leg stimulation [77], communicating a state of a variable via wrist movements [45]).

Unlike these interactive systems that use EMS as a form of force feedback or as an information channel, we explore EMS in a new direction: leveraging user's unique muscular responses to EMS as a form of active biometric authentication.

2.2 Biometric authentication

Biometric authentication verifies an individual by their unique biological characteristics. To verify a user's identity, a biometric authentication system compares a previously stored biometric key from a particular user to incoming, typically real-time, biometric data of the user wishing to authenticate. Compared to traditional password or PIN based methods, biometric authentication offers significantly better usability by not requiring the user to memorize passwords or PINs.

Existing biometric systems can be categorized into two types: passive and active biometrics.

Passive biometrics. Passive biometrics rely on physiological characteristics that naturally occur in users, which can be either static or dynamic. Static data, e.g., fingerprints [27], handprints [22], facial and eye features [2, 51, 59, 85], is often used for authentication. Biometrics based on dynamic data recognize patterns that vary over time, e.g., heartbeats [31], gait [78], mouse movements [32], keystrokes [80], speech features [4], body movements [54, 62], pulse-response [67] and bioimpedance [28, 70]. Compared to static data, these display greater complexity and are harder to model.

Passive biometrics are vulnerable to data thefts and replay attacks as reported by numerous incidents and studies [6, 18, 35, 57, 86–88]. This is because the identity (also known as "key") associated with each user is physically "hard-coded" and then used *repeatedly* for all authentications. Thus after a key has been compromised (e.g., stolen from a database), an adversary can bypass authentication until the key is replaced. Finally, there is a small number of available biological traits per user that act as suitable keys, e.g., once all ten fingerprints are compromised, this user can never again rely on fingerprint authentication.

Active biometrics via challenge-response. Active biometrics leverage a user's physiological response to a given stimulus (also known as "challenge") injected by the interactive device. The assumption is that each user's response to a given challenge is unique. Thus, each *challenge-response* is effectively a biometric password. Examples of challenge-response biometrics include leveraging: the palm's response to vibrations [40], reflexive eye behaviors in response to visual stimuli [73], or even EEG responses [42]. These systems authenticate implicitly so the user does not need to consciously follow the challenge, e.g., the palm vibrates and the user is authenticated [40].

Compared to passive biometrics, active systems are more robust against data thefts and replay attacks. This is because each user can potentially generate many challenges, each triggering a different response. The system uses a new challenge in each authentication session, preventing attackers from using previously observed responses to breach it.

Lastly, while many challenge-response authentication systems leverage the user's movement (e.g., gaze [66] or wrist shakes [60]), these require explicit action from the user. Unlike these, our novel exploration of EMS-based authentication provides the advantages of movement-based challenge-response while automatically delivering the challenge and eliciting the user's *involuntary* response.

3 IMPLEMENTATION

To help readers replicate our design, we provide the necessary technical details. Furthermore, to accelerate replication, we provide source code and training scripts¹. Here, we describe in detail the prototype we implemented for our user studies, which is based on sensing the user's movements using inertial measurement units (IMUs). However, this is just one possible configuration for our concept. As depicted in Figure 1, other tracking systems, such as optical tracking [56, 84], are likely feasible alternatives.

3.1 System Overview

ElectricAuth consists of three components: (1) a medically-compliant **EMS device** that delivers EMS challenges to the user,

¹http://sandlab.cs.uchicago.edu/electricauth

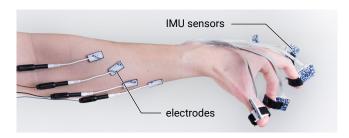


Figure 2: IMU-based version of our EMS authentication system, which we used for our user studies.

(2) a **motion sensor** that captures the actuated limb's movements, such as IMUs or depth cameras, and (3) a **trained machine learning model** that classifies the user's movements and performs authentication. Figure 2 depicts one concrete implementation of our system using EMS and IMUs attached to a user's forearm, which we used for our user studies.

1. EMS hardware.

EMS stimulator: For delivering EMS impulses we use the Hasomed Rehastim, a medical compliant device with eight individually controllable channels. This device has often been used in interactive systems based on EMS [47–49]. To control the EMS stimulation, our software sends serial commands via USB using the Hasomed's Science Protocol [24]. These impulses have a latency of <1ms.

Customized EMS sleeve: As with any device based on EMS, we start by calibrating the electrode placement for each user at her registration session. Our calibration aims at targeting four muscles on the user's forearm that actuate finger and wrist rotation. At the anterior forearm we stimulate two muscle groups: (1) primarily the flexor carpi radialis and partially the flexor digitorum profundus; and, (2) the flexor pollicus longus. At the posterior forearm we stimulate two muscle groups: (1) primarily the extensor digitorum and partially the extensor digiti minimi, extensor pollicis brevis & longus; and, (2) the extensor indicis. As is typical with EMS-based systems, these electrode positions are adjusted for each user during the registration session to ensure comfort. Because each user has a different muscular anatomy and body shape, the resulting electrode locations are different across different users.

After calibration, the resulting electrode layout for a particular user is fixed by making an EMS-electrode sleeve (fabric with electrodes stitched to it) that this user wears any time they use ElectricAuth. Moreover, the sleeve becomes part of each user's own challenge definition, i.e., an attacker trying to impersonate a particular user will require obtaining or copying the user's sleeve, which we later validate in our studies by actually providing the impersonators with the EMS sleeves of the legitimate users.

EMS parameters: Our EMS stimuli on all electrode locations are the same: single-shot square-impulses with an intensity of 10mA and a pulse-width of $200\mu s$. We chose this configuration for two reasons. First, we configured EMS impulses to generate small and subtle finger movements rather than large conspicuous movements typical of most existing EMS research, because this enables more practical authentication scenarios. While these smaller movements are harder to recognize, our results suggest that our authentication model can accurately track these (see Section 7). Second, we opted to

make all impulses uniform to shine light in the fact that intersubject variability in EMS arises from factors external to EMS waveform characteristics.

Our EMS challenges are constructed by sequencing these standardized pulses to one of the four channels the user's forearm is connected to. For instance, one can construct a challenge with a sequence of six impulses, each followed by a resting period. We detail the engineering of our pulse sequences in Section 3.2.

2. Motion sensing.

We utilized a set of five 9-DOF inertial measurement units, attached to the fingers via a 3D printed ring (NXP Precision 9DoF, comprised of the FXOS8700 3-Axis accelerometer and magnetometer and the FXAS21002 3-axis gyroscope). These sample the fingers' acceleration and rotation at 50Hz (post-sample interpolated to 100Hz) with a precision of $\pm 4g$ at 14-bit for acceleration and $\pm 250^\circ/s$ at 16-bit for rotations; note that we do not use the magnetometer. These IMUs are sampled by a ATSAMD21G18 ARM Cortex M0 48 MHz processor, via a TCA9548A I2C Multiplexer. Finally, our sensing board relays the IMU data via serial over USB to our software.

While attaching IMUs to each finger has been shown to be a reliable way to capture hand pose [15, 29], we believe many alternative tracking systems are possible to realize EMS-based authentication, such as depth cameras [72, 84], RGB cameras [9, 71], and others [34]. We provide a short evaluation that confirmed the use of depth cameras as an alternative tracking system in Section 9.

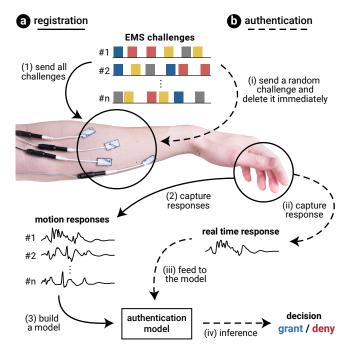


Figure 3: Interactive pipeline for the registration (registering a new user) and authentication phase (interactive use in runtime). User response can be captured using a motion capturing device, e.g., IMUs and cameras (not shown). In this system, the EMS device and electrodes are wearable; the motion capturing device is either wearable or placed near the user; while the authentication model can be remote.

3. Authentication software and pipeline.

The software component of ElectricAuth, written in Python, handles all the interactions between EMS device, motion sensing, model training and real-time authentication. The pipeline of ElectricAuth, which is depicted in Figure 3, is comprised of two phases: (a) **registration** and (b) **authentication**.

In the **registration** phase, marked by solid lines in Figure 3, registering a new user (after calibration) is as follows: (1) a set of *n* EMS challenges are sent one at a time; (2) the user's movements in response to each challenge are recorded; (3) these responses are used to train a machine learning-based authentication model for this user. The number of challenge-response recorded per user is the primary factor that dictates the total time the system needs for registering a single user (we detail this in Section 4).

In the **authentication** phase, marked by dashed lines in Figure 3, verifying a user's identity in run-time is as follows: (i) one random EMS challenge belong to the claimed identity is chosen, deleted immediately from the database, and sent to the user via EMS; (ii) the user's movements in response to the challenge are recorded; (iii) the motion responses are fed into the trained authentication model of the claimed identity; (iv) the system determines whether this user is legitimate (i.e., being the claimed identity) or not.

3.2 Engineering EMS-based Challenges

As our system is the first that explores EMS for authentication, we dedicated a significant part of our exploration in understanding how to increase the challenge pool using EMS; a large challenge pool is what makes a challenge-response based authentication system robust against data breaches and replay attacks. Naively, one can stimulate the user's muscles with individually configurable pulses; however, this (1) requires more calibration time and (2) does not reveal the mechanisms that explain these individual responses. Therefore, we kept purposely all EMS impulses uniform for all users of our system; this grants us more confidence in interpreting the unique responses as originating from the physiological differences between users. Yet, this introduces a challenge when it comes to diversifying the challenge pool.

One straightforward solution (adopted by many existing works on challenge-response biometrics [40, 42, 73]) is to sequence stimuli but separate them by a fixed time gap. If we were to adopt this as well, the maximum number of EMS challenges would be S^L , where S is the number of unique stimuli in the system and L is the number of stimuli in each challenge. For example, a sequence of six EMS impulses over four possible EMS channels, with a fixed rest period between each impulse, results in $4^6 = 4,096$ challenges. We were interested in whether we could dramatically surpass this approach.

To significantly increase our challenge pool, we explored a rather unused property of human muscles that causes them to respond differently to EMS depending on their current state of contraction. We call this *temporal dependency*.

Temporal dependency. We empirically found, in our preliminary pilots, that a subject's response to an EMS stimulus is affected by the previous stimulus in the same challenge, and the impact depends on the time gap between them (represented as τ).

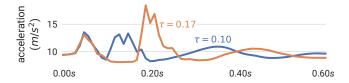


Figure 4: An example of how a response changes when the time gap between two EMS stimuli varies: we vary the time gap from 0.1s (blue curve) to 0.17s (orange curve).

Figure 4 shows two example traces of a finger's acceleration when we stimulate the user's muscles with a sequence of two stimuli (A and B) but vary the time gap between A and B (i.e., $\tau=0.1s$ and $\tau=0.17s$). The measured acceleration displays different characteristics when we vary τ . The strongest candidate for a physiological explanation is that muscle contractility and elasticity vary with muscle length [21, 81], and the response to a stimulus depends on the muscle lengths at the time of stimulation. Thus, depending on the gap between A and B, the subject's unique contractility [23] and elasticity [26, 82] will lead to different responses.

The use of temporal dependency affords a large EMS challenge pool by varying the time gaps between consecutive stimuli. Assuming they all produce distinguishable responses, the number of unique challenges (of length L) is upper bounded by $S^L \cdot T^{L-1}$, where T is the number of distinct time gaps. For our ElectricAuth prototype, we utilize S=4 EMS channels and T=7 different time gaps ($\tau=\frac{1}{30}$ s, $\frac{2}{30}$ s, ..., $\frac{7}{30}$ s), which in early pilots we found to lead to sufficiently different movement outcomes. The maximum number of unique challenges is 112 (L=2), 87, 808 (L=4) or 68, 841, 472 (68M) (L=6), compared to 16 (L=2), 256 (L=4) or 4, 096 (L=6) when we do not vary the time gap.

Further increasing the challenge pool. Encoding longer challenges is another way to expand the challenge pool. With S=4 stimulus locations and T=7 time gaps, sending L=8 pulses (<2s) increases the pool size to 53, 971, 714, 048 ($4^8 \times 7^7$). Also it is possible to add more electrodes or customize EMS impulses to further diversify the pool.

Checking for uniqueness. Ideally, every challenge-response authentication in the pool is unique. However, in practice this might not be the case given the granularity and sensitivity of motion sensors. To enforce uniqueness, ElectricAuth can apply a verification step during user registration. Specifically, after generating new challenges for a user at the registration phase, it collects the corresponding responses and checks the similarity across these responses and previously registered responses (e.g., computing the mean square error (MSE) between raw responses). If a new challenge is identified as a previously registered challenge, this new challenge is removed.

4 USER AUTHENTICATION MODEL

We now present the design of ElectricAuth's user authentication model. ElectricAuth requires a trained authentication model per legitimate user, which is used to verify whether a test subject is indeed that user. To do so, the model takes as input the response to a given challenge designed for the legitimate user, and outputs whether the test subject is legitimate. Our authentication model

was designed with two objectives in mind: (1) minimize the amount of samples collected from the user (i.e., reducing registration overhead) and (2) resist common attacks (e.g., impersonation and replay attacks) and data breaches.

4.1 Overview

Initially, we explored implementing our model using specific features of the user's IMU data in response to particular EMS challenges (so called feature-engineering). However, we quickly realized a major downside of this approach: as the response data we capture in real-time from the IMUs is complex (thirty concurrent data streams: 5×3 axes of acceleration and 5×3 axes of rotation), simple feature extraction might not capture the full expressivity of the data. Therefore, after experimenting with this approach, we turned to neural network based models.

We implemented a robust authentication model that, for each registered user, integrates two deep neural network (DNN) models to resist both impersonation and replay attacks. Specifically, authentication starts with (1) an unsupervised anomaly detector, which verifies whether a response was produced by the user the model belongs to (i.e., the legitimate user); this step prevents *impersonation attacks*, in which a different user attempts to gain identity of the legitimate user. If a response passes the anomaly detector, it then enters (2) a challenge classifier, which detects and rejects *replay attacks* by verifying whether the response is the reaction to the challenge used in the current authentication session.

Both models are trained using only the challenge-response pairs of this legitimate user collected during registration. When the user (re)registers a new set of challenges, we retrain both models from scratch using the new data. This also enables ElectricAuth to recover from data and model breaches.

4.2 Detailed Model Design

1. Verifying user via unsupervised anomaly detection.

We implement user verification as unsupervised anomaly detection [7], where the detection model is trained on *only* the legitimate user's responses collected during registration. At run time, the model verifies whether an input response was likely originated by the legitimate user. This anomaly-based detector leverages the fact that responses from other users will display characteristics different from those of the legitimate user. Thus the model is designed to produce normal output when the input response comes from its legitimate user, but abnormal output when the input comes from any other user. This design prioritizes generality as the model is trained *without* requiring knowledge on other users.

For our prototype, we apply a reconstruction error based anomaly detection system [63]. Specifically, we use variational autoencoder (VAE) [16], a DNN architecture well-known for *automatically* capturing complex patterns in target data. As shown in Figure 5, each VAE starts from an encoder to extract latent features from each input response, followed by a decoder to reconstruct the response from these features. It then computes the mean squared error (MSE) between the input and reconstructed responses, and outputs it as the anomaly score of the input. Ideally, the anomaly score will be low when the input response comes from the legitimate user

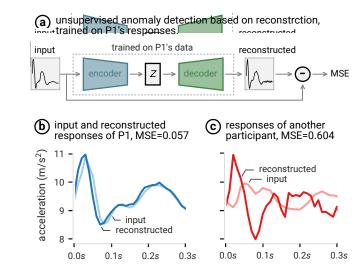


Figure 5: Authentication starts with an anomaly detection, which verifies if a response came from the legitimate user that the model belongs to (P1 in this example). (a) The anomaly score is the MSE of the input and model-reconstructed responses. We illustrate how our anomaly detector correctly: (b) identifies P1 (legitimate user) with a low MSE and (c) rejects P2 (impersonator) with a high MSE.

and high when the input comes from a different user. Thus, the system can configure a threshold on the anomaly score, where a value larger than the threshold indicates the test subject is not the legitimate user (i.e., the user verification fails). In ElectricAuth, we choose the threshold during model training to reach a desired false rejection rate (i.e., the probability that the model rejects the legitimate user's input responses).

For our implementation, we train our VAE using each legitimate user's responses to all the chosen challenges collected during registration. The data aggregation (across challenges) creates a reasonable amount of data to train the VAE successfully. We consider a common VAE architecture [19], where the encoder contains two dense layers of 400 and 200 neurons, respectively, and the decoder contains two dense layers of 400 and 3600 neurons, respectively, to match the input size.

To illustrate the effectiveness of our model, we plot in Figure 5b-c the input and reconstructed responses of a legitimate user (here, P1 of our user study) and a different participant P2 (also from our user study), respectively, using the model trained for P1. For the sake of visual clarity, we only plot the responses from only one accelerometer axis. Both responses are not used for model training. We see that P1's response is well-approximated by the model-reconstructed response; in fact, with a very low MSE of 0.057. On the other hand, P2's response (when tested on P1's model), produces a large MSE of 0.604, around 10-fold higher than the MSE of the legitimate user (P1).

Figure 6 shows the responses (collected by the IMUs) of five subjects (P1-5) to a challenge designed for P1 (the legitimate user in this case). When tested on P1's anomaly detection model, the



Figure 6: Sample responses of a P1's challenge (with L=6 impulses) and impersonators' responses (P2, P3, P4 and P5) to the same challenge. Each row is a sensor channel and each column is one data sample. Here we show one second of responses. When tested on P1's anomaly detection model, the corresponding anomaly scores for P1-5 are 0.70, 5.03, 9.44, 8.81 and 7.50, respectively. In this case, the model can easily detect impersonators.

anomaly scores for these responses are 0.70, 5.03, 9.44, 8.81 and 7.50, respectively. Thus P1's model easily rejects P2-5 as impersonators.

2. Verifying challenge via challenge classifier.

Next in the authentication pipeline, ElectricAuth verifies whether the input response matches the challenge used in the current session. As mentioned earlier, this is designed to resist replay attacks, where an attacker, after obtaining a copy of the legitimate user's responses to previously used challenges, replays one of these responses to bypass authentication.

ElectricAuth implements challenge verification by training a classifier: given an input response, it determines the corresponding challenge. If the identified challenge matches the challenge used in the current authentication session, authentication is granted; otherwise, rejected. Moreover, the classifier also detects when the input response comes from any challenge *not* used to train the classifier, because the classifier will output a low confidence score.

Our implementation uses a Convolutional Neural Network (CNN) for this classification task [58]. It contains four convolutional and two dense layers. Each convolutional layer employs 64 filters sized 5 to extract information from the input. The information is then fed into the two dense layers containing 128 and 112 neurons, respectively. At the end, a softmax function is applied to the output to produce a probability distribution over potential challenges. We train our CNN using the *same* training data used in training the above anomaly detector, except that we now label each response by its corresponding challenge.

5 CONTRIBUTIONS, BENEFITS AND LIMITATIONS

Our main contribution is that we explore EMS in a new direction, i.e., leveraging EMS's *intersubject variability* as a novel modality for active biometric authentication.

ElectricAuth inherits the advantages of both biometric and password authentication: (1) As with any biometric authentication device, ElectricAuth does not require memorization or cognitive effort – this makes our system suited for a wide range of users, including those with cognitive impairments; (2) Unlike passive biometrics (such as fingerprints), ElectricAuth's challenge-response structure makes it secure against data breaches and replay attacks; Lastly,

(3) ElectricAuth leverages temporal dependency to create a very large set of challenges – in this way, ElectricAuth can dispose a challenge anytime like a one-time password.

On the flipside, ElectricAuth is subject to several limitations: (1) Like any solution based on electrical muscle stimulation, ElectricAuth requires some initial adjustments of the electrodes (during registration) that ensure pain-free operation, and also periodic regelling of adhesive electrodes to prevent electrodes from fatigue and eventually affecting the authentication accuracy; (2) ElectricAuth requires user's hands to be free while authenticating, making it more suitable for hands-free applications; (3) As with existing biometric devices, ElectricAuth requires initial registration. Specifically, each challenge needs to be registered in advance; Lastly, (4) while a single EMS impulse can be very short (e.g., 200μ s) to achieve very high accuracy, we expanded our sequence to 1.2 seconds of muscle stimulation, as such ElectricAuth takes ~1300ms to authenticate a user in runtime. While this is certainly fast enough for most applications, it is longer than some passive approaches, such as fingerprint recognition.

6 OVERVIEW OF EVALUATIONS

We evaluated our concept of using EMS for authentication by means of four different evaluations, each shining light on a different facet of our research question. All studies were approved by our Institutional Review Board (IRB no. omitted for anonymity). To aid the reader in understanding the different validations we performed, we present an overview of our evaluations with a preview of their respective results:

I. User studies. We evaluated the feasibility of EMS as an active biometric with three experiments and 13 participants. We found that that ElectricAuth resists three common attacks: (1) impersonations attacks, in which participants played impersonators against each legitimate user (attack success rate or false acceptance rate: 0.17%); (2) replay attacks, in which participants mimic the movements of the legitimate user from videos (success rate: 0.00%), or replay a perfect record of response to any used challenges directly into the IMUs (success rate: 0.00%); and, (3) synthesis attacks, in which we synthesized data from the participants' data to attack their authentication models (success rate: 0.2-2.5%).

II. Exploratory longitudinal study. We conducted a longitudinal study over 24 days and for two participants, to examine ElectricAuth's authentication model over time and against various muscle conditions (fatigue, humidity, etc.). We found that an authentication model, trained using the first three days and tested over the next 21 days, performed very stable over time and on muscle conditions unseen during training (false rejection rate $\approx 2\%$, with a SD around 3%)

III. Technical evaluation. A technical in which we measured ElectricAuth's latency, model training time, and the feasibility of using depth cameras as an alternative motion tracking modality.

IV. Testing model robustness at scale, using synthetic data. We applied a data-driven approach to better understand how our system might scale to larger numbers of users that is simply impractical to test in the laboratory. To realize this, we employed the user study data to train deep generative models that produce synthetic impersonator responses, and used these data to further evaluate ElectricAuth. We found that, across all the data-driven experiments and for all legitimate users, no generated response was accepted by ElectricAuth (attack success rate: 0).

7 USER STUDIES

To evaluate the feasibility of EMS as an active biometric we conducted a user study, with three sub-experiments, which allowed us to understand: (1) authentication accuracy, in which we evaluated the accuracy of our system; (2) impersonation attack, in which we evaluated its robustness against attackers trying to impersonate legitimate users; and, (3) replay attack and synthesis attack, in which we evaluated its robustness against three replay attacks (human mimicry, record-replay, breach-replay) and one online synthesis attack.

In total, we collected 70,000+ wrist and finger movements as responses to EMS challenges (stimulation patterns). We analyzed the performance of ElectricAuth using four standard metrics, typically employed to assess a system's authentication performance: (1) False rejection rate (FRR), which measures how often a legitimate user is mistakenly denied, at a specific threshold; (2) False acceptance rate (FAR), which measures how often an illegitimate user is mistakenly authorized, at a specific threshold; (3) Equal error rate (EER), the rate at which the measured FRR equals the measured FAR for a certain threshold; and, (4) Receiver operator characteristic curve (ROC curve), which describes the relationship between FRR and FAR as a curve, by varying its threshold.

7.1 Experiment#1: Authentication Accuracy

The goal of our first study was to understand the authentication accuracy of our system. Furthermore, as we were interested in the impact of the length of the EMS challenges on its performance, we recorded participants' movements to three sets of challenges, based on their number of impulses L=1,2,6 (referred to as length-1, -2, and -6 challenges, respectively). For each challenge set we stimulated participants' forearms and recorded finger movements using IMUs.

Participants. We recruited 13 participants from our institution (mean age= 24 years, SD= 3 years; mean weight= 66.3 kg, SD=

13.3 kg; mean height= 171.2 cm, SD= 8.2 cm; 7 females, 6 males). Participants were compensated with 50 USD for their time.

Apparatus. Participants wore our system on their left forearm. This included the EMS and IMU components, which were fitted by an experimenter. To ensure participants' comfort with EMS, we calibrated it so that all electrode channels operated pain-free. To ensure that all target muscles were correctly stimulated (see Implementation for details), we gradually increased the intensity during calibration, following calibration process similar to [5]. If a participant felt any discomfort before reaching the target intensity, we moved to another electrode position. To minimize fatigue, participants rested their elbow on a resting base.

After calibration, we recorded each participant's exact electrode locations by making a custom sleeve with marked positions. These 13 sleeves were later used in Experiment #2, where we examined impersonation attacks (i.e., each impersonator wore the sleeve of a legitimate user to attack our authentication system).

During the study, participants did not receive any specific instruction, since we wanted them to react naturally to the EMS impulses.

EMS challenges. The EMS challenges in our study were configured as previously described, i.e., a challenge was comprised of a sequence of single-shot square-impulses with an intensity of 10mA and a pulse-width of $200\mu s$; these sequences were of length-1, -2 or -6. In between each pair of impulses we included a time gap. Each gap was one of seven possible durations $(\frac{1}{30}s, \frac{2}{30}s, ..., \frac{6}{30}s, \frac{7}{30}s)$; thus, the recording duration of a length-1, -2, and -6 challenges were 0.6s, 0.8s and 1.2s, respectively. While length-1 challenges were collected in this experiment, these were only used for an analysis in Experiment#2 (anomaly detector performance).

Procedure. To test whether ElectricAuth correctly authenticates our 13 participants, we first registered each participant. Our system did this automatically: (1) a participant feels an EMS challenge, (2) their forearm muscles react involuntarily, and (3) our system records the response. We repeated this process 10 times per challenge. These ten responses were shuffled to remove potential sequence effects. These responses were then randomly divided into a training set (eight responses) and a testing set (two responses). Then, our system took these eight responses (for all challenges) and trained the anomaly detector and challenge classifier for each participant. As cross-validation, we repeated this process to produce 10 authentication models per participant and reported the average test results of these models in all our subsequent experiments.

For length-1 and -2 challenges, we tested the full set of challenges (a total of four for length-1 and 112 for length-2). For length-6 challenges, we were forced to test only a subset, since the full set includes 68, 841, 472 challenges, which would be fatiguing for participants. Therefore, we randomly chose 115 challenges from the full set.

In total, each participant performed 2310 trials: 40 trials of the four length-1 challenges (10 repetitions); 1120 trials of the 112 length-2 challenges (10 repetitions); and, 1150 trials of the 115 length-6 challenges subset (10 repetitions).

Results: overall authentication accuracy. We first examine the accuracy of the end-to-end authentication model, which depends on the accuracy of both the anomaly detection model and

	planned FRR			P7	2.5	5.0
participant	2%	5%		P8	2.3	5.5
P1	2.3	6.1	_	P9	3.2	5.8
P2	2.1	5.5		P10	2.2	5.0
P3	2.1	4.9		P11	2.3	5.0
P4	1.7	5.4		P12	2.8	5.5
P5	2.6	4.8		P13	2.6	5.4
P6	2.7	6.2	_	AVG(SD)	2.4(0.4)	5.4(0.4)

Table 1: The measured false rejection rate (FRR, %) for all registered participants (P1-P13) closely matched the planned FRR. The measured FRR was calculated for each participant using their test responses to 115 length-6 challenges.

the challenge classification model. We defined overall accuracy as the probability that a legitimate response successfully passed the two-step authentication. Note that the accuracy is dependent on the anomaly threshold used by ElectricAuth's authentication model. During model training, we configured the threshold to reach a planned false rejection rate (FRR). Note that the threshold is determined using just the training data (without the knowledge of any run-time testing data). Ideally, the run-time measured FRR (i.e., 1–accuracy) should equal to the planned FRR.

For each of the 13 registered participants, Table 1 summarizes the measured FRR (i.e., = 1–accuracy) aggregating the results across all 115 challenges (of length 6). Here we reported the results for planned FRR of 2% and 5%. We see that the measured FRR closely matched the planned FRR. Across all the participants, the mean measured FRR is 2.4% (SD of 0.4%) and 5.4% (SD of 0.4%), respectively, matching the two planned FRR values (2% and 5%).

Results: challenge classification accuracy. Digging deeper into the accuracy of our system, we turn to evaluate the accuracy of challenge classification model (as it is the main component protecting against replay attacks). Our accuracy findings are depicted in Figure 7. For length-2 challenges (complete set, i.e., 112 of them) the average accuracy is 99.89% (SD=0.19% across users). And for length-6 subset of challenges we found an accuracy of 99.78% (SD=0.50%). These results also show that the challenges (full set of length-2, subset of length-6) are unique across each other.



Figure 7: Electric Auth's challenge classification accuracy for length-2 and length-6 challenges.

7.2 Experiment#2: Impersonation Attacks

In this user study, we measured our system's ability to resist impersonation attacks.

Participants. For this study, we invited all 13 participants from Study#1. Participants were briefed that they would play an attacker trying to impersonate other participants. Participants were compensated with 50 USD for their time.

Procedure & apparatus. For each target participant, we applied their customized challenges (used in Experiment #1) to the other 12 participants (as impersonators) and collected their responses. Impersonators were asked to wear the sleeve fabricated for each target participant in Experiment #1. These sleeves grant the impersonator with the exact electrode positions of the legitimate user. We also tested cases where impersonators wear their own sleeves and other participants' sleeves and found that wearing the target participant's sleeve leads to the most effective attack; thus we focused on it.

In total, each participant performed 3240 trials: 480 trials of the length-1 challenges (10 repetitions per challenge, impersonated 12 other participants); 2760 trials of the length-6 challenges subset (2 repetitions per challenge, impersonated 12 other participants).

Impersonating someone else by using their electrode placement does not guarantee comfortable use, i.e., we did not adjust electrodes to preserve the legitimate participant's placement. While no participant felt uncomfortable with length-1 challenges, there was some discomfort on a few length-6 trials (3.8% of the total); anytime a participant voiced discomfort, we stopped the stimulation and discarded this trial.

Results: performance of anomaly detector. To deepen our understanding of intersubject variability and the anomaly detection model performance, we first compared the responses to a single stimulus (or length-1 challenge), submitted by each target participant in Experiment#1 and the 12 impersonators in this experiment. We fed these responses to the target participant's anomaly detection model and recorded their anomaly scores. For the sake of visual clarity, we normalized these anomaly score values by the target participant's average anomaly score value (see Figure 8).

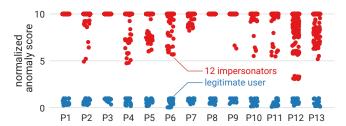
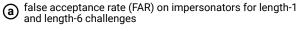


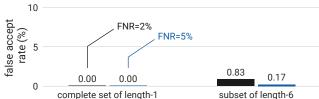
Figure 8: Normalized reconstruction error for the responses to each participant's length-1 challenges, submitted by both the legitimate user and the 12 impersonators. For visual clarity, we capped the value at 10.

We found that our anomaly detector for each participant is well-trained and can distinguish impersonators from the legitimate participant. This is clear as Figure 8 depicts a large separation between the legitimate participant and the impersonators. It also confirms EMS intersubject variability.

Results: robustness against impersonation attack. We examined the end-to-end success rate of impersonation attacks against each participant, using the attack data collected on length-1 challenges (complete set) and length-6 challenges (the 115 subset).

Figure 9(a) depicts the false acceptance rate (aggregated across 13 participants' models since they are consistent) against length-1 and length-6 challenges, for the planned FRR of 2% and 5%, respectively. With length-1 challenges (4 challenges), the impersonation attack





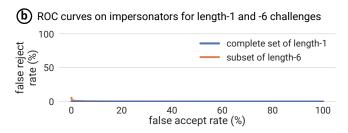


Figure 9: ElectricAuth's robustness against impersonation attacks.

failed. With length-6 challenges, the attack exhibited a very low success rate, only 0.83% (SD=1.14%) at planned FRR=2% and 0.17% (SD=0.32%) at planned FRR=5%. Again this suggests that our system is robust against impersonation attacks. Figure 9(b) shows the ROC curves under impersonation attacks with length-6 challenges, where ElectricAuth achieves an EER of 1.31%.

7.3 Experiment#3: replay and synthesis attacks

In this user study, we measured ElectricAuth's robustness against replay attacks and synthesis attacks, both trying to engineer a response to bypass authentication after obtaining some knowledge on the legitimate user's responses.

We considered three replay attacks, and one synthesis attack, ranging in increased attack complexity:

- (1) human mimicry, where the attacker video-tapes and studies a participant's responses and then physically mimics the responses without wearing any EMS;
- **(2) record-replay**, where the attacker compromises the IMUs so that they can *perfectly* record the target participant's response to challenges in previous authentication sessions, then during a new authentication session (i.e., a new challenge), the attacker selects a previous recorded response and directly feeds it to the IMUs;
- (3) breach-replay, where the attacker breaches the database or the model to recover stored challenge-response data, and feeds one response to the IMU's circuit; here ElectricAuth reacts to the breach by asking users to re-register using new challenges and retraining the models:
- (4) online synthesis, where the attacker compromises both EMS and IMUs to record both the challenge and the response in previous sessions; then at run-time, the attacker searches through these records and attempts to synthesize and submit in *real-time* an engineered response to the current challenge. For these attacks, we evaluated ElectricAuth using the false acceptance rate (FAR) and the ROC curve.

Participants. We recruited five participants to perform the human mimicry attack: three from our previous study (chosen at random) and two new participants from our local institution (ages: 25 & 22 years old; weights: $55 \& 99 \ kg$; heights: $177 \& 180 \ cm$; one female and one male). Participants were compensated with $50 \ \text{USD}$ for their time.

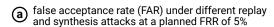
Procedure. In the **human mimicry attack**, we asked participants to study 23 videos of finger movements of a target participant. Each video was a recording of one single response to a length-6 challenge. Participants were allowed to study these videos as many times as they intended and in slow-motion (recorded at 240 fps, with clear and unobstructed view of the finger movements). Once confident and ready, participants were asked to mimic these finger movements while wearing only the IMU component of our system, in their best attempt to impersonate the target participant. Furthermore, as reference, we also asked the target participant that had partaken in Experiment#1 to self-mimic 23 of his own EMS responses after observing and studying them.

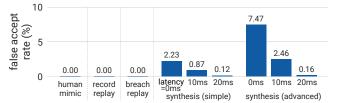
Results: robustness against human mimicry. We found that none of the study participants was able to fool our system by mimicking the target participant's responses. Note that these participants were allowed to view the videos in slow motion and as many times as they want. The FAR was 0 for a FRR $\geq 2\%$. This confirms our intuition that the EMS movements are indeed involuntary and incredibly hard to voluntarily replicate.

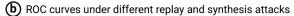
Results: robustness against record-replay attack. For this we utilized data from Experiment#1. Even assuming perfect recording on the side of the attacker (i.e., their recording channel has access to IMUs without any noise or sample rate issues), we found our system to be robust against these attacks. In particular, for length-6 challenges, the FAR (against any of the 115 challenges) was less than 0.0014% across all 13 participants when FRR \geq 2%. This FAR is significantly smaller than the challenge misclassification rate of our authentication model (0.2%, see Experiment#1).

Results: robustness against breach-replay attack. Again we utilized data from Experiment#1. For each participant, we randomly split the 115 challenges (and their responses) into two equal sets (A and B). We assume that the attacker, via data breach, obtains the dataset A and uses them to launch replay attacks against ElectricAuth. At the same time, ElectricAuth reacts to the data breach by asking users to re-register via a set of new challenges (i.e., dataset B) and retraining the authentication models using dataset B. Like the above, we found our system to be robust against these replay attacks − the FAR was less than 0.0098% when FRR ≥ 2%. Moreover, both the anomaly detector and challenge classifier components in the model were able to reject the attack responses.

Results: robustness against online synthesis. We evaluated the success rate of an online synthesis attack, using the data from Experiment#1. We assume the attacker has access to the EMS and IMUs without sample rate or noise issues, which is in itself very unlikely. The idea behind a synthesis attack is that the adversary records both challenges and their responses, and segments these into chunks, as in "this impulse at electrode 1, moves this finger by this much", and so forth. We referred to this approach as the simple synthesis attack. A more advanced attack would capture the







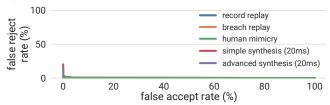


Figure 10: ElectricAuth's robustness against different replay and synthesis attacks. For online synthesis, the attacker had perfect records on responses to 50 challenges. Here ElectricAuth operates on length-6 challenges.

impact of temporal dependency by segmenting responses into per pair-stimuli chunks, as in "these two impulses at electrodes 1 and 2, move these fingers by this much"). After segmenting the responses, the attacker will observe each incoming impulse of a new challenge and inject a response into the IMUs in real-time. Note that even assuming best hardware and knowledge, assembling this response will always have some latency.

Figure 10(a) plots the FAR of online synthesis attacks considering three latency values, assuming the attacker has observed R=50 challenge-response pairs and the planned FRR is set to 5%. Even under the extreme attack case (zero latency, which is physically impossible), the attack success rate is low (i.e., FAR=2.2% and 7.5% for simple and advanced attacks, respectively). When the synthesis latency reaches 20ms, which still depicts an unlikely extremely fast response, the FAR drops to 0.1-0.2%. The same applies when we raised R to 75 (i.e., the advanced attack's success rate is only 0.25% for latency=20ms).

Results: ROC and EER. Finally, Figure 10(b) plots the ROC results for all the replay attacks and synthesis attacks (with latency =20ms). We see that ElectricAuth achieves noticeable EERs only for the synthesis attacks (1.48% for simple synthesis and 1.59% for advanced synthesis). These results show that ElectricAuth is robust against replay and synthesis attacks, even those extreme ones.

8 EXPLORATORY LONGITUDINAL STUDY

We conducted an exploratory longitudinal study to examine ElectricAuth over time and against various environment and muscle conditions. Specifically, we performed **fixed-model-over-time tests**, which depicts how an authentication model trained using the first three days of data will perform over time and under muscle conditions (e.g., humidity, fatigue, etc.) and other non-predictable environmental factors that were not present in the training data;

Participants. Due to Covid-19, only two co-authors participated in this study (ages: 25 & 24 years old; weights: 70 & 54kg; heights 170 & 163cm; one male and one female).

Procedure. In the day prior to the start of the 24-day period, we conducted an initial calibration session (following the same method and apparatus described in Experiment #1). Then, we followed with 24 days of data collection. We collected data once a day. For each participant, we randomly chose 115 length-6 challenges to collect user responses.

For this study, we used fabric sleeves with embedded EMS electrodes at the precalibrated positions for each participant, following a design similar to [36]. Each day, participants were asked to wear their custom electrode-sleeves (depicting their calibrated locations). Participants fitted the sleeve by themselves prior to the trials by aligning markings on the sleeve with their elbow and top of wrist. If the electrode pads were dry, they re-gelled it using conductive gel. Then, they recorded their response to the 115 challenges every day. For each challenge, they collected more than 6 responses per day. After the trials, they removed the sleeve until the next day.

Conditions. To explore the impact of environmental and physiological variations, we conducted data collection under combinations of three conditions: (1) time of the day (morning/afternoon/late); (2) environment humidity (dry/damp); and (3) muscle fatigue (normal/fatigued). We randomly chose one combination per day, and each combination was tested at least twice during the study. In the damp condition, participants were asked to stay in their bathroom with the humidity at over 80% and temperature over 29 °C for more than 20 minutes right before the data collection. For dry condition, participants stayed in an air-conditioned room of humidity 55% and temperature 24°C. To test our system right after the muscles started to fatigue, participants were asked to do a routine of intense forearm muscle training (dumbbell wrist flexion and extension) for a minimum of 15 minutes before collecting data.

During the days in which we tested ElectricAuth under normal muscle conditions, participants still performed their forearm muscle training but after the data collection session. This allowed us to study if extended muscle exercise would affect the system performance.

Training the authentication model. For each participant, we used data collected in the first three days to train the authentication model (the anomaly detector and challenge classifier). For both participants, the training data were collected under the same (dry, pre-workout) condition. The rest of the data (21 days) were used for testing our authentication models. The testing data contained conditions both seen or unseen in the training data. We excluded day 10 and 11 for participant 1 due to need for replenishing the sticky gel on the electrodes, i.e., waiting for gel supply.

For all the trained models, we configured the anomaly detection thresholds to achieve a planned FRR of 2%. As discussed before, such configuration is set using only the training data without the knowledge of any testing data.

Results: fixed-model-over-time tests. To understand the impact of a specific condition (time of the day, environment humidity or muscle fatigue), Fig. 11(a)(b) shows the measured false rejection rate (FRR) under each condition. For both participants, the measured FRRs are reasonably consistent across conditions and

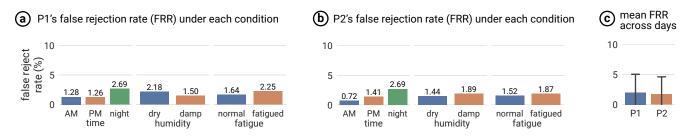


Figure 11: Results of fixed-model-over-time tests. (a) and (b) shows for both participants, our system is stable under various conditions; (c) our system is stable over time (21 days) for both participants.

closely match the planned FRR (2%). But more importantly, while our authentication models are trained only under the combination of dry and pre-workout conditions, they remain accurate under other conditions not seen during training. This provides initial evidence on the generality of ElectricAuth.

For both participants, we also plot the mean FRRs over time in Fig. 11(c). We see that the FRR is stable over time, (mean=2.01%, SD=3.13%) for participant 1 and (mean = 1.76%, SD=2.90%) for participant 2. No significant performance degradation over time was found for both participants. These results suggest that ElectricAuth remains relatively stable on a monthly scale.

9 TECHNICAL EVALUATION

We deepened our understanding of how future interactive systems might be built based on EMS authentication by measuring system latency, training time, and the feasibility of depth cameras as an alternative tracking modality.

9.1 Authentication latency

To measure ElectricAuth's inference latency (i.e., time needed to make a decision in run-time) and the model training, we utilized the data from the participants of Experiment#1, i.e., 115 length-6 challenges, eight response records per challenge for training, two response records for testing.

Run-time inference latency. As we probed the future of EMS-based authentication, we were interested in understanding how ElectricAuth would perform on smaller platforms, such as laptops or even embedded devices. As such, we ran our system on a MacBook Pro with a Intel Core i9-9880H CPU and on a Nvidia Jetson Nano embedded device (measuring 70 x 45 mm). Our results show that our system can authenticate a user in 3ms on laptop's CPU and 35ms on a small embedded device. This result suggests our approach is feasible for quick authentications and even available on mobile or wearable devices.

Training latency. Our results demonstrate that it took 35s (33s for anomaly detector; 2s for the challenge classifier) to train the complete model on a Nvidia Titan RTX GPU and 542s on a laptop's CPU (501s for anomaly detector; 41s for the challenge classifier).

9.2 Using camera to capture finger movements

While we used IMUs to capture finger movements in our user study, we believe these movements can also be captured via other modalities, such as depth cameras, a common platform for hand pose estimation [72, 84]. To test our belief, we carried out a simple feasibility experiment. Here, we swapped out IMU sensors with a RGB-D camera (Intel RealSense D435), which operates at 640x480 resolution and 30 frames per second. The camera was placed in front of the participant with a distance of 50cm.

Following the same procedure of Experiment#1, we recorded, via the depth camera, the responses to our 115 length-6 challenges on one participant. We then used an available hand gesture recognition model (from [38]) as our challenge verification model.

We found that the challenge classification accuracy for this simple feasibility experiment was 99.57% using the depth image. We also measure a 0.00% success rate of a record-replay attack against this participant's model.

10 USING SYNTHETIC DATA TO TEST ATTACKS AT SCALE

Our user study demonstrated that ElectricAuth was accurate in verifying each of the 13 participants and robust against any attacks in that scale. However, gaining insight into how ElectricAuth would perform in larger deployments (e.g., 100's of users) is impractical by means of user studies at an early stage. To shed light into this, we explore a data-driven approach to evaluate ElectricAuth's robustness against impersonation attacks using synthetic data.

Procedure. We followed the recent approach of generating synthetic data by training deep generative models, which is shown to produce diverse and natural data (e.g., objects [69], human faces [3, 89], faces with emotions [50], and physiological data including ECG, EEG, and so forth [25]) beyond the training set. Specifically, we used the PixelCNN++ model [69], a state-of-the-art deep generative model for images (since we treat each response as an image). Following [69], we trained a generative model for each legitimate user in our experiment #2 (see Section 7.2), using the impersonator responses collected for this user (12 subjects and 115 challenges), conditioned on the challenge. Once trained, the generator produces random, natural variations of the training data, emulating responses of potential impersonators beyond our user study. We validated each generator using the well-known negative log likelihood (NLL), which produced results on par with (and often slightly better than) those reported by [69] on object/face images. This indicated that our trained generators are able to learn and follow the actual data distribution rather than overfitting to the training data.

Results: robustness against synthetic impersonators. For each of the 13 users in our experiment #2, we used the corresponding generator to produce 1075 impersonator responses against this user. These include 100 synthetic impersonators for each of 5 randomly selected challenges, and 5 additional impersonators for each of 115 challenges. We then tested these impersonator responses on ElectricAuth's authentication model for this user (i.e., the same authentication model used in our experiment #2). All impersonator responses were rejected (i.e., 0% FAR at 5% FRR). This result aligns with our user study results, and sheds lights on ElectricAuth's robustness against impersonation attacks at larger scales.

11 CONCLUSIONS, APPLICATIONS & FUTURE WORK

We proposed, implemented and evaluated the use of electrical muscle stimulation (EMS) as a novel modality for active biometrics. We engineered an interactive system, which we called ElectricAuth, that stimulates the user's forearm muscles with a sequence of electrical impulses (i.e., an EMS challenge) and measures the user's involuntary finger movements (i.e., response to the challenge). The key idea behind ElectricAuth is that it leveraged EMS's intersubject variability, i.e., the same electrical stimulation results in different movements in different users because everybody's physiology is unique (e.g., differences in bone and muscular structure, skin resistance and composition, etc.). Moreover, we demonstrated that ElectricAuth is secure against data breaches and replay attacks, as it never reuses the same challenge twice in authentications - the key property that allowed ElectricAuth to achieve this is that in just one second of stimulation our system was able to encode one of 68M possible challenges.

11.1 Potential applications

We believe that ElectricAuth is applicable to a range of interactive scenarios in which users authenticate without needing to memorize passwords or PINs. We believe this is of special interest for devices that natively offer motion tracking or finger tracking, such as for virtual reality (which we illustrated in Figure 1 using the Oculus Quest), smartwatch-based interaction [52, 83, 90] or even leveraging a smartphone's built in IMUs. Furthermore, we believe our approach is of particular interest for accessibility scenarios, such as authentication for users with motor-impairments (e.g., spinal cord injury, arguably the most impactful application of EMS in the medical domain [61]) but with intact musculature.

11.2 Future work

We believe this first exploration of EMS for user authentication provides fertile grounds for exploring subsequent challenges and opportunities: (1) while we have shown ElectricAuth worked well on the full set of 112 length-2 challenges and a subset of 115 length-6 challenges, growing the size of a challenge might enable new applications, as such, research is needed to demonstrate that this approach works across an even larger set of challenges and over a longer time period; (2) while ElectricAuth worked well on the 13 participants from our user studies, more physiological research is needed to deepen understanding of EMS's intersubject variability;

(3) while ElectricAuth worked well on the controlled wrist posture, more investigation is required to understand its performance under other postures and their impacts; lastly, (4) as new EMS systems emerge from the medical domain (e.g., higher resolution electrode arrays [33, 36, 65], implanted devices [64], and so forth), a system like ElectricAuth will likely improve in wearability and performance, which will require further investigations.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive feedback. We also thank Michael Maire for his suggestions on synthetic data generation, and Yuxin Guan for her help on data analysis. This work is supported in part by the National Science Foundation grant CNS-1949650. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agencies.

REFERENCES

- Ali Ahanger and Anil Kumar. 2014. Effect of Anthropometric Factors on Motor Nerve Conduction Velovity in Healthy Kashmiri Population. *International Journal of Medical and Applied Sciences* 3 (feb 2014), 125–132.
- [2] Karan Ahuja, Rahul Islam, Varun Parashar, Kuntal Dey, Chris Harrison, and Mayank Goel. 2018. EyeSpyVR: Interactive Eye Sensing Using Off-the-Shelf, Smartphone-Based VR Headsets. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 2, 2, Article 57 (July 2018), 10 pages. https://doi.org/10.1145/3214260
- [3] A. Ali-Gombe, E. Elyan, and C. Jayne. 2019. Multiple Fake Classes GAN for Data Augmentation in Face Image Dataset. In 2019 International Joint Conference on Neural Networks (IJCNN). 1–8.
- [4] Aware. 2020. Voice Authentication. https://www.aware.com/voice-authentication/.
- [5] Xueliang Bao, Yuxuan Zhou, Yunlong Wang, Jianjun Zhang, Xiaoying Lü, and Zhigong Wang. 2018. Electrode placement on the forearm for selective stimulation of finger extension/flexion. *PloS one* 13, 1 (2018).
- [6] Thomas Brewster. 2019. We Broke Into A Bunch Of Android Phones With A 3D-Printed Head. https://www.forbes.com/sites/thomasbrewster/ 2018/12/13/we-broke-into-a-bunch-of-android-phones-with-a-3d-printedhead/#4a9a79213307.
- [7] Saikiran Bulusu, Bhavya Kailkhura, Bo Li, Pramod K. Varshney, and Dawn Song. 2020. Anomalous Instance Detection in Deep Learning: A Survey. arXiv:2003.06979 [cs.LG]
- [8] Joseph P Campbell. 1997. Speaker recognition: A tutorial. Proc. IEEE 85, 9 (1997), 1437–1462.
- [9] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. 2019. OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019).
- [10] R. Chandler, C. Clauser, J. McConville, Herbert Reynolds, and J. Young. 1975. Investigation of Inertial Properties of the Human Body. (03 1975), 171.
- [11] Charles E Clauser, John T McConville, and J W Young. 1971. Weight, Volumn, and Center of Mass of Segments of The Human Body. Journal of Occupational and Environmental Medicine (1971).
- [12] Cory Cornelius, Ronald Peterson, Joseph Skinner, Ryan Halter, and David Kotz. 2014. A Wearable System That Knows Who Wears It. In Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (Bretton Woods, New Hampshire, USA) (MobiSys '14). Association for Computing Machinery, New York, NY, USA, 55–67. https://doi.org/10.1145/2594368.2594369
- [13] A Desplantez, C Cornu, and F Goubel. 1999. Viscous properties of human muscle during contraction. *Journal of biomechanics* 32, 6 (1999), 555–562. https://doi. org/10.1016/S0021-9290(99)00039-1
- [14] K.N. Dewangan, G.V. [Prasanna Kumar], P.L. Suja, and M.D. Choudhury. 2005. Anthropometric dimensions of farm youth of the north eastern region of India. *International Journal of Industrial Ergonomics* 35, 11 (2005), 979 – 989. https://doi.org/10.1016/j.ergon.2005.04.003
- [15] Laura Dipietro, Angelo M Sabatini, and Paolo Dario. 2008. A survey of glove-based systems and their applications. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews) 38, 4 (2008), 461–482.
- [16] Carl Doersch. 2016. Tutorial on Variational Autoencoders arXiv:1606.05908 [stat.ML]
- [17] R Drillis, R Contini, and M Bluestein. 1964. Body Segment Parameters; A Survey of Measurement Techniques. Artificial limbs 8 (1964), 44–66. http://europepmc. org/abstract/MED/14208177

- [18] Simon Eberz, Nicola Paoletti, Marc Roeschlin, Andrea Patané, Marta Kwiatkowska, and Ivan Martinovic. 2017. Broken Hearted: How To Attack ECG Biometrics. In 24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February 26 March 1, 2017. The Internet Society. https://www.ndss-symposium.org/ndss2017/ndss-2017-programme/broken-hearted-how-attack-ecg-biometrics/
- [19] Basic VAE Example. 2019. https://github.com/pytorch/examples/tree/master/vae.
- [20] Farzam Farbiz, Zhou Hao Yu, Corey Manders, and Waqas Ahmad. 2007. An Electrical Muscle Stimulation Haptic Feedback for Mixed Reality Tennis Game. In ACM SIGGRAPH 2007 Posters (San Diego, California) (SIGGRAPH '07). Association for Computing Machinery, New York, NY, USA, 140–es. https://doi.org/10.1145/ 1280720.1280873
- [21] Michael J. Grey. 1997. Viscoelastic properties of the human wrist during the stabilization phase of a targeted movement. Master's thesis. Simon Fraser University, Burnaby, Canada.
- [22] Anhong Guo, Robert Xiao, and Chris Harrison. 2015. CapAuth: Identifying and Differentiating User Handprints on Commodity Capacitive Touchscreens. In Proceedings of the 2015 International Conference on Interactive Tabletops Surfaces (Madeira, Portugal) (ITS '15). Association for Computing Machinery, New York, NY, USA, 59–62. https://doi.org/10.1145/2817721.2817722
- [23] SDR Harridge, R Bottinelli, M Canepari, MA Pellegrino, C Reggiani, M Esbjörnsson, and B Saltin. 1996. Whole-muscle and single-fibre contractile properties and myosin heavy chain isoforms in humans. Pflügers Archiv 432, 5 (1996), 913–920.
- [24] Hasomed. 2020. https://hasomed.de/en/.
- [25] Andres Hernandez-Matamorosb, Hamido Fujita, and Hector Perez-Meanac. 2020. A novel approach to create synthetic biomedical signals using BiRNN. Elsevier Information Sciences 541 (Dec 2020), 218–241.
- [26] AL Hof. 1998. In vivo measurement of the series elasticity release curve of human triceps surae muscle. *Journal of biomechanics* 31, 9 (1998), 793–800.
- [27] Christian Holz and Patrick Baudisch. 2013. Fiberio: A Touchscreen That Senses Fingerprints. In Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (St. Andrews, Scotland, United Kingdom) (UIST '13). Association for Computing Machinery, New York, NY, USA, 41–50. https://doi. org/10.1145/2501988.2502021
- [28] Christian Holz and Marius Knaust. 2015. Biometric Touch Sensing: Seamlessly Augmenting Each Touch with Continuous Authentication. In Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (Charlotte, NC, USA) (UIST '15). Association for Computing Machinery, New York, NY, USA, 303–312. https://doi.org/10.1145/2807442.2807458
- [29] Christopher-Eyk Hrabia, Katrin Wolf, and Mathias Wilhelm. 2013. Whole Hand Modeling Using 8 Wearable Sensors: Biomechanics for Hand Pose Prediction. In Proceedings of the 4th Augmented Human International Conference (Stuttgart, Germany) (AH '13). Association for Computing Machinery, New York, NY, USA, 21–28. https://doi.org/10.1145/2459236.2459241
- [30] Report: Data Breach in Biometric Security Platform Affecting Millions of Users. 2019. https://www.vpnmentor.com/blog/report-biostar2-leak/.
- [31] Steven A Israel and John M Irvine. 2012. Heartbeat biometrics: a sensing system perspective. *International Journal of Cognitive Biometrics* 1, 1 (2012), 39–65.
- [32] Zach Jorgensen and Ting Yu. 2011. On mouse dynamics as a behavioral biometric for authentication. In Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security. 476–482.
- [33] T. Keller, M. Lawrence, A. Kuhn, and M. Morari. 2006. New Multi-Channel Transcutaneous Electrical Stimulation Technology for Rehabilitation. In 2006 International Conference of the IEEE Engineering in Medicine and Biology Society. 194–197.
- [34] David Kim, Otmar Hilliges, Shahram Izadi, Alex D. Butler, Jiawen Chen, Iason Oikonomidis, and Patrick Olivier. 2012. Digits: Freehand 3D Interactions Anywhere Using a Wrist-Worn Gloveless Sensor. In Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST'12). Association for Computing Machinery, New York, NY, USA, 167–176. https://doi.org/10.1145/2380116.2380139
- [35] Tomi Kinnunen, Md. Sahidullah, Héctor Delgado, Massimiliano Todisco, Nicholas Evans, Junichi Yamagishi, and Kong Aik Lee. 2017. The ASVspoof 2017 Challenge: Assessing the Limits of Replay Spoofing Attack Detection. In *Proc. Interspeech* 2017. 2–6. https://doi.org/10.21437/Interspeech.2017-1111
- [36] Jarrod Knibbe, Paul Strohmeier, Sebastian Boring, and Kasper Hornbæk. 2017. Automatic Calibration of High Density Electric Muscle Stimulation. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 1, 3, Article 68 (Sept. 2017), 17 pages. https://doi.org/10.1145/3130933
- [37] Michinari Kono, Takumi Takahashi, Hiromi Nakamura, Takashi Miyaki, and Jun Rekimoto. 2018. Design Guideline for Developing Safe Systems That Apply Electricity to the Human Body. ACM Trans. Comput.-Hum. Interact. 25, 3, Article 19 (June 2018), 36 pages. https://doi.org/10.1145/3184743
- [38] Okan Köpüklü, Ahmet Gunduz, Neslihan Kose, and Gerhard Rigoll. 2019. Realtime Hand Gesture Detection and Classification Using Convolutional Neural Networks. In 14th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2019, Lille, France, May 14-18, 2019. IEEE, 1-8. https://doi.org/10.

- 1109/FG.2019.8756576
- [39] Ernst Kruijff, Dieter Schmalstieg, and Steffi Beckhaus. 2006. Using Neuromuscular Electrical Stimulation for Pseudo-Haptic Feedback. In Proceedings of the ACM Symposium on Virtual Reality Software and Technology (Limassol, Cyprus) (VRST '06). Association for Computing Machinery, New York, NY, USA, 316–319. https: //doi.org/10.1145/1180495.1180558
- [40] Jingjie Li, Kassem Fawaz, and Younghyun Kim. 2019. Velody: Nonlinear Vibration Challenge-Response for Resilient User Authentication. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (London, United Kingdom) (CCS '19). Association for Computing Machinery, New York, NY, USA, 1201–1213. https://doi.org/10.1145/3319535.3354242
- [41] Chong L. Lim, Chris Rennie, Robert J. Barry, Homayoun Bahramali, Ilario Lazzaro, Barry Manor, and Evian Gordon. 1997. Decomposing skin conductance into tonic and phasic components. *International Journal of Psychophysiology* 25, 2 (1997), 97 – 109. https://doi.org/10.1016/S0167-8760(96)00713-1
- [42] Feng Lin, Kun Woo Cho, Chen Song, Wenyao Xu, and Zhanpeng Jin. 2018. Brain Password: A Secure and Truly Cancelable Brain Biometrics for Smart Headwear. In Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services (Munich, Germany) (MobiSys '18). Association for Computing Machinery, New York, NY, USA, 296–309. https://doi.org/10.1145/ 3210240.3210344
- [43] Pedro Lopes and Patrick Baudisch. 2013. Muscle-Propelled Force Feedback: Bringing Force Feedback to Mobile Devices. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Paris, France) (CHI '13). Association for Computing Machinery, New York, NY, USA, 2577–2580. https://doi.org/10.1145/2470654.2481355
- [44] Pedro Lopes, Alexandra Ion, and Patrick Baudisch. 2015. Impacto: Simulating Physical Impact by Combining Tactile Stimulation with Electrical Muscle Stimulation. In Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (Charlotte, NC, USA) (UIST '15). Association for Computing Machinery, New York, NY, USA, 11–19. https://doi.org/10.1145/2807442.2807443
- [45] Pedro Lopes, Alexandra Ion, Willi Mueller, Daniel Hoffmann, Patrik Jonell, and Patrick Baudisch. 2015. Proprioceptive Interaction. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 939–948. https://doi.org/10.1145/2702123.2702461
- [46] Pedro Lopes, Patrik Jonell, and Patrick Baudisch. 2015. Affordance++: Allowing Objects to Communicate Dynamic Use. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 2515–2524. https://doi.org/10.1145/2702123.2702128
- [47] Pedro Lopes, Sijing You, Lung-Pan Cheng, Sebastian Marwecki, and Patrick Baudisch. 2017. Providing Haptics to Walls & Heavy Objects in Virtual Reality by Means of Electrical Muscle Stimulation. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 1471–1482. https: //doi.org/10.1145/3025453.3025600
- [48] Pedro Lopes, Sijing You, Alexandra Ion, and Patrick Baudisch. 2018. Adding Force Feedback to Mixed Reality Experiences and Games Using Electrical Muscle Stimulation. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, Article 446, 13 pages. https://doi.org/10.1145/ 3173574.3174020
- [49] Pedro Lopes, Doăa Yüksel, François Guimbretière, and Patrick Baudisch. 2016. Muscle-Plotter: An Interactive System Based on Electrical Muscle Stimulation That Produces Spatial Output. In Proceedings of the 29th Annual Symposium on User Interface Software and Technology (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 207–217. https://doi.org/10.1145/ 2984511.2984530
- [50] Y. Luo and B. Lu. 2018. EEG Data Augmentation for Emotion Recognition Using a Conditional Wasserstein GAN. In 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). 2535–2538.
- [51] Davide Maltoni, Dario Maio, Anil K Jain, and Salil Prabhakar. 2009. Handbook of fingerprint recognition. Springer Science & Business Media.
- [52] Meethu Malu, Pramod Chundury, and Leah Findlater. 2018. Exploring Accessible Smartwatch Interactions for People with Upper Body Motor Impairments. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18). Association for Computing Machinery, New York, NY, USA, 12. https://doi.org/10.1145/3173574.3174062
- [53] Wayne J Millar. 1986. Distribution of body weight and height: comparison of estimates based on self-reported and observed measures. Journal of Epidemiology & Community Health 40, 4 (1986), 319–323.
- [54] Tahrima Mustafa, Richard Matovu, Abdul Serwadda, and Nicholas Muirhead. 2018. Unsure How to Authenticate on Your VR Headset? Come on, Use Your Head!. In Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics (Tempe, AZ, USA) (IWSPA '18). Association for Computing Machinery, New York, NY, USA, 23–30. https://doi.org/10.1145/3180450

- [55] Gergana Stefanova Nikolova and Yuli Emilov Toshev. 2007. Estimation of male and female body segment parameters of the Bulgarian population using a 16segmental mathematical model. *Journal of Biomechanics* 40, 16 (2007), 3700–3707. https://doi.org/10.1016/j.jbiomech.2007.06.016
- [56] Oculus. 2019. Introducing Hand Tracking on Oculus Quest-Bringing Your Real Hands into VR. https://www.oculus.com/blog/introducing-hand-tracking-onoculus-quest-bringing-your-real-hands-into-vr/.
- [57] Patrick Howell O'Neill. 2019. Data leak exposes unchangeable biometric data of over 1 million people. https://www.technologyreview.com/f/614163/data-leakexposes-unchangeable-biometric-data-of-over-1-million-people/.
- [58] Francisco Javier Ordóñez and Daniel Roggen. 2016. Deep Convolutional and LSTM Recurrent Neural Networks for Multimodal Wearable Activity Recognition. Sensors (Basel) 16, 1 (Jan 2016). https://doi.org/10.3390/s16010115
- [59] Marcos Ortega, M.G. Penedo, J. Rouco, N. Barreira, and M.J. Carreira. 2009. Personal verification based on extraction and characterisation of retinal feature points. *Journal of Visual Languages & Computing* 20, 2 (2009), 80–90. https://doi.org/10.1016/j.jvlc.2009.01.006
- [60] Shwetak N. Patel, Jeffrey S. Pierce, and Gregory D. Abowd. 2004. A Gesture-Based Authentication Scheme for Untrusted Public Terminals. In Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology (Santa Fe, NM, USA) (UIST '04). Association for Computing Machinery, New York, NY, USA, 157–160. https://doi.org/10.1145/1029632.1029658
- [61] P. Hunter Peckham and Jayme S. Knutson. 2005. Functional Electrical Stimulation for Neuromuscular Applications. *Annual Review of Biomedical Engineering* 7, 1 (2005), 327–360. https://doi.org/10.1146/annurev.bioeng.6.040803.140103
- [62] Ken Pfeuffer, Matthias J. Geiger, Sarah Prange, Lukas Mecke, Daniel Buschek, and Florian Alt. 2019. Behavioural Biometrics in VR: Identifying People from Body Motion and Relations in Virtual Reality. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, Article 110, 12 pages. https://doi.org/10.1145/3290605.3300340
- [63] Marco A.F. Pimentel, David A. Clifton, Lei Clifton, and Lionel Tarassenko. 2014. A review of novelty detection. Signal Processing 99 (2014), 215 – 249. https://doi.org/10.1016/j.sigpro.2013.12.026
- [64] D. Popovic, L. L. Baker, and G. E. Loeb. 2007. Recruitment and Comfort of BION Implanted Electrical Stimulation: Implications for FES Applications. IEEE Transactions on Neural Systems and Rehabilitation Engineering 15, 4 (2007), 577– 586.
- [65] Ana Popović-Bijelić, Goran Bijelić, Nikola Jorgovanović, Dubravka Bojanić, Mirjana B. Popović, and Dejan B. Popović. 2005. Multi-Field Surface Electrode for Selective Electrical Stimulation. Artificial Organs 29, 6 (2005), 448–452. https://doi.org/10.1111/j.1525-1594.2005.29075.x
- [66] Vijay Rajanna, Seth Polsley, Paul Taele, and Tracy Hammond. 2017. A Gaze Gesture-Based User Authentication System to Counter Shoulder-Surfing Attacks. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI EA '17). Association for Computing Machinery, New York, NY, USA, 1978–1986. https://doi.org/10.1145/ 3027063.3053070
- [67] Kasper Bonne Rasmussen, Marc Roeschlin, Ivan Martinovic, and Gene Tsudik. 2014. Authentication Using Pulse-Response Biometrics. In 21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23-26, 2014. The Internet Society. https://www.ndss-symposium. org/ndss2014/authentication-using-pulse-response-biometrics
- [68] Brian Reed. 1997. The physiology of neuromuscular electrical stimulation. Pediatric Physical Therapy 9, 3 (1997), 96–102. https://journals.lww.com/pedpt/Fulltext/1997/00930/The_Physiology_of_Neuromuscular_Electrical.2.aspx
- [69] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma. 2017. Pixel-CNN++: A PixelCNN Implementation with Discretized Logistic Mixture Likelihood and Other Modifications. In ICLR.
- [70] Munehiko Sato, Rohan S. Puri, Alex Olwal, Yosuke Ushigome, Lukas Franciszkiewicz, Deepak Chandra, Ivan Poupyrev, and Ramesh Raskar. 2017. Zensei: Embedded, Multi-Electrode Bioimpedance Sensing for Implicit, Ubiquitous User Recognition. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3972–3985. https://doi.org/10.1145/3025453. 3025536
- [71] T. Simon, H. Joo, I. Matthews, and Y. Sheikh. 2017. Hand Keypoint Detection in Single Images Using Multiview Bootstrapping. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 4645–4653. https://doi.org/10. 1109/CVPR.2017.494
- [72] Ayan Sinha, Chiho Choi, and Karthik Ramani. 2016. DeepHand: Robust Hand Pose Estimation by Completing a Matrix Imputed With Deep Features. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 4150–4158.

- https://doi.org/10.1109/CVPR.2016.450
- [73] Ivo Sluganovic, Marc Roeschlin, Kasper B. Rasmussen, and Ivan Martinovic. 2016. Using Reflexive Eye Movements for Fast Challenge-Response Authentication. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (Vienna, Austria) (CCS '16). Association for Computing Machinery, New York, NY, USA, 1056–1067. https://doi.org/10.1145/2976749.2978311
- [74] Spiceworks. 2018. Spiceworks Study Reveals Nearly 90 Percent of Businesses Will Use Biometric Authentication Technology by 2020. https://www.spiceworks.com/press/releases/spiceworks-study-reveals-nearly-90-percent-businesses-will-use-biometric-authentication-technology-2020/.
- [75] Diana S. Stetson, James W. Albers, Barbara A. Silverstein, and Robert A. Wolfe. 1992. Effects of age, sex, and anthropometric factors on nerve conduction measures. *Muscle & Nerve* 15, 10 (1992), 1095–1104. https://doi.org/10.1002/mus. 880151007
- [76] Primož Strojnik, Alojz Kralj, and I Ursic. 1979. Programmed six-channel electrical stimulator for complex stimulation of leg muscles during walking. IEEE Transactions on Biomedical Engineering 2 (1979), 112–116.
- [77] Fangmin Sun, Chenfei Mao, Xiaomao Fan, and Ye Li. 2019. Accelerometer-Based Speed-Adaptive Gait Authentication Method for Wearable IoT Devices. *IEEE Internet of Things Journal* 6, 1 (2019), 820–830. https://doi.org/10.1109/JIOT.2018. 2860592
- [78] Fangmin Sun, Chenfei Mao, Xiaomao Fan, and Ye Li. 2019. Accelerometer-Based Speed-Adaptive Gait Authentication Method for Wearable IoT Devices. *IEEE Internet of Things Journal* 6, 1 (2019), 820–830. https://doi.org/10.1109/JIOT.2018. 2860592
- [79] Emi Tamaki, Takashi Miyaki, and Jun Rekimoto. 2011. PossessedHand: Techniques for Controlling Human Hands Using Electrical Muscles Stimuli. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Vancouver, BC, Canada) (CHI '11). Association for Computing Machinery, New York, NY, USA, 543–552. https://doi.org/10.1145/1978942.1979018
- [80] Pin Shen Teh, Andrew Beng Jin Teoh, and Shigang Yue. 2013. A survey of keystroke dynamics biometrics. The Scientific World Journal 2013 (2013). https://doi.org/10.1155/2013/408280
- [81] The Muscle Physiology Laboratory UCSD. 2008. Fundamental Functional Properties of Skeletal Muscle. http://muscle.ucsd.edu/musintro/props.shtml.
- [82] Kai Uffmann, Stefan Maderwald, Waleed Ajaj, Craig G Galban, Serban Mateiescu, Harald H Quick, and Mark E Ladd. 2004. In vivo elasticity measurements of extremity skeletal muscle with MR elastography. NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In Vivo 17, 4 (2004), 181–190.
- [83] Tran Huy Vu, Archan Misra, Quentin Roy, Kenny Choo Tsu Wei, and Youngki Lee. 2018. Smartwatch-Based Early Gesture Detection 8 Trajectory Tracking for Interactive Gesture-Driven Applications. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 2, 1 (March 2018), 27. https://doi.org/10.1145/3191771
- [84] Chengde Wan, Thomas Probst, Luc Van Gool, and Angela Yao. 2019. Self-Supervised 3D Hand Pose Estimation Through Training by Fitting. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 10845–10854. https://doi.org/10.1109/CVPR.2019.01111
- [85] Richard P Wildes. 1997. Iris recognition: an emerging biometric technology. Proc. IEEE 85, 9 (1997), 1348–1363.
- [86] Davey Winder. 2019. Apple's iPhone FaceID Hacked In Less Than 120 Seconds. https://www.forbes.com/sites/daveywinder/2019/08/10/apples-iphone-faceid-hacked-in-less-than-120-seconds/#449ff4b621bc.
- [87] Davey Winder. 2019. Hackers Claim Any Smartphone Fingerprint Lock Can Be Broken In 20 Minutes. https://www.forbes.com/sites/daveywinder/2019/11/02/ smartphone-security-alert-as-hackers-claim-any-fingerprint-lock-broken-in-20-minutes/#588a90116853.
- [88] Yi Xu, True Price, Jan-Michael Frahm, and Fabian Monrose. 2016. Virtual U: Defeating Face Liveness Detection by Building Virtual Models from Your Public Photos. In 25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016, Thorsten Holz and Stefan Savage (Eds.). USENIX Association, 497-512.
- [89] Jian Zhao, Lin Xiong, Karlekar Jayashree, Jianshu Li, Fang Zhao, Zhecan Wang, Sugiri Pranata, Shengmei Shen, Shuicheng Yan, and Jiashi Feng. 2017. Dual-Agent GANs for Photorealistic and Identity Preserving Profile Face Synthesis. In Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 65–75.
- [90] Junhan Zhou, Yang Zhang, Gierad Laput, and Chris Harrison. 2016. AuraSense: Enabling Expressive Around-Smartwatch Interactions with Electric Field Sensing. In Proceedings of the 29th Annual Symposium on User Interface Software and Technology (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 81–86. https://doi.org/10.1145/2984511.2984568