Treadmill Assisted Gait Spoofing (TAGS): An Emerging Threat to Wearable Sensor-based Gait Authentication

RAJESH KUMAR, Syracuse University and Haverford College
CAN ISIK and VIR V. PHOHA, Syracuse University

In this work, we examine the impact of Treadmill Assisted Gait Spoofing on Wearable Sensor-based Gait Authentication (WSGait). We consider more realistic implementation and deployment scenarios than the previous study, which focused only on the accelerometer sensor and a fixed set of features. Specifically, we consider the situations in which the implementation of WSGait could be using one or more sensors embedded into modern smartphones. In addition, it could be using different sets of features or different classification algorithms, or both. Despite the use of a variety of sensors, feature sets (ranked by mutual information), and six different classification algorithms, Treadmill Assisted Gait Spoofing was able to increase the average false accept rate from 4% to 26%. Such a considerable increase in the average false accept rate, especially under the stringent implementation and deployment scenarios considered in this study, calls for a further investigation into the design of evaluations of WSGait before its deployment for public use.

CCS Concepts: • **Security and privacy** → **Authentication; Biometrics**;

Additional Key Words and Phrases: User authentication, behavioral biometrics, gait authentication, gait spoofing, wearable sensors

ACM Reference format:

Rajesh Kumar, Can Isik, and Vir V. Phoha. 2021. Treadmill Assisted Gait Spoofing (TAGS): An Emerging Threat to Wearable Sensor-based Gait Authentication. *Digit. Threat.: Res. Pract.* 2, 3, Article 23 (June 2021), 17 pages.
<https://doi.org/10.1145/3442151>

1 BACKGROUND AND MOTIVATION

More than 100 studies have been published over the past two decades on **Wearable Sensor-based Gait Authentication (WSGait)** [9, 33, 44, 51, 54, 61]. Most of these studies have evaluated the uniqueness of WSGait under the zero-effort attack scenario. Only a handful of studies have paid attention to circumvention of WSGait via imitation (does not include data injection) [16–19, 27, 36, 38, 55]. To the best of our knowledge, all but two [27, 55] of these studies have shown that it is non-trivial to circumvent WSGait via imitation. With a limited experimental setup, Kumar et al. [27] showed that one could spoof WSGait with comparative ease by adjusting gait characteristics such as speed, step length, step width, and thigh lift with the help of an off-the-shelf treadmill. The scope of the previous work [27] is limited to the accelerometer sensor and a fixed set of features.

Authors' addresses: R. Kumar, Syracuse University and 370 Lancaster Ave, Haverford College, PA, 19041; email: rkuma102@syr.edu; C. Isik, 131 Link Hall, Syracuse University, NY, 13210; email: cisik@syr.edu; V. V. Phoha, Syracuse University, 4-183 CST, Syracuse University, NY, 13210; email: vvphoha@syr.edu.

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.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2576-5337/2021/06-ART23 \$15.00

<https://doi.org/10.1145/3442151>

Furthermore, the same set of features was used to train the imitator, which could have favored the attack. Therefore, the motivation is to explore the applicability of treadmill attack under more stringent scenarios, including those in which WSGait would be using readings of one or more sensors embedded into the device and a different feature set than the one used for training the imitator.

2 RELATED WORK

This work primarily extends the idea presented in Kumar et al. [27]. The notable differences between this and the previous work [27] are as follows:

- The authentication systems tested for circumvention in Kumar et al. [27] were built using only one sensor's readings (i.e., accelerometer readings). In practice, however, the authentication system could be using readings of multiple sensors embedded in smart devices [31, 32]. Previous studies have reported that the fusion of multiple sensors helps achieve better error rates [28, 31, 32]. Thus, it becomes a logical next step to test the effectiveness of the attack presented in earlier work [27] on systems that use the readings of more than one sensor for user authentication. Since it is not always possible to find out the sensors in use, we decided to study four commonly used sensors (accelerometer, gyroscope, magnetometer, and rotation vectors) and all of their possible combinations [22, 28, 31, 32].
- The attack presented in Kumar et al. [27] uses the same set of features for both training the imitator, and training and evaluating the authentication models. In other words, although unstated, the Kumar et al. [27] assume that the attacker would know the set of features used to implement the authentication system subject to the attack. In practice, this assumption may not be valid. The knowledge of the feature space makes the model even more susceptible to attack. For example, Zhao et al. [63] showed that the knowledge of the length of the feature space helps launch a data injection attack on WSGait. Therefore, it is likely that each user authentication model is implemented using a different set of features selected by a feature selection algorithm (see Figure 1). Therefore, a logical extension of the attack presented in earlier work [27] is to test its effectiveness on the authentication systems that may or may not be using the same set of features that the attacker used to train the imitator.
- We elaborate on the working of the attack, in what scenarios it applies, how to access the required information needed to launch the attack and improve its success, advantages over other attack methods, limitations, and possible countermeasures.

Besides Kumar et al. [27], the closely related works are listed in the Table 1. We can see that the implementation strategies of mimicry attacks have evolved over the past decade. Gafurov et al. [16–19] investigated the security of gait under two types of deliberate spoofing scenarios: friendly and hostile. Under the friendly scenario, the imitators did not make any deliberate attempt to copy the genuine user, whereas under the hostile scenario, individuals with similar physical characteristics attempted to imitate each other. The chances of impostor acceptance did not increase in either of the scenarios. The authors concluded that the imitator with similar physical characteristics or of the same gender might have a better chance of getting accepted [16].

In a follow-up, Stang [55] analyzed five templates of the same user collected at different settings of gait factors such as speed and step length. The author employed a total of 13 imitators; each made 15 attempts to match each template. The whole experiment took about 30 minutes. The imitators had not seen the target walking before they started the imitation process. Live plots of resultant accelerations were shown to the imitators on a big screen in addition to the match scores (Pearson's correlation coefficient). The match score was computed between the patterns of the imitator and the target user. Some imitators were able to exceed the correlation coefficients of 0.5 (considered a 50% match) criteria set by the authors, but randomly which led to the conclusion that one may imitate the gait patterns of targeted individuals if trained rigorously. Later, Mjaaland et al. [36] pointed out that

Table 1. List of the Most Closely Related Work, Publication Year, Number of Subjects, Gait Cycles per Subject, and Imitators per Subject, Along with the Number of Cycles and Error Under the Zero- and Mimicry-effort Attacks

Authors	Year	Users (cycles/user)	Imitators per User (cycles)	Avg. Error (%) (before)	Avg. Error (%) (after)	Remarks
Gafurov et al. [19]	2006	22 (15)	1 (2)	16	16	No substantial threat
Gafurov et al. [16]	2007	100 (4)	1 (4)	13	13	Same gender a threat
Gafurov et al. [18]	2007	100 (4)	1 (4)	13	13	Same gender or closest in dataset is a threat
Stang [55]	2007	1 (5)	1 (13)	26	–	Correlation for assessing success, only one target user
Mjaaland et al. [36]	2011	50 (10)	–	6.2	6.2	Training makes it worse
Kumar et al. [27]	2015	18 (532)	1 (970)	FAR: 6	SFAR: 38	Controlled training on a treadmill
Muaaz & Mayrhofer [38]	2017	35 (200)	5 victims & 4 imitators	13	13	Training makes it worse, used actors
This work	2020	18 (532)	1 (970)	FAR: 4	SFAR: 26	Multiple sensors, Frame based, and Dynamically selected features for each authentication model

Note: The size of the data is provided in terms of number of gait cycles to make a comparison among different works easy. In this work, however, we used a sliding window-based segmentation approach that output fixed-length frames (see Section 3.1.2). Each of the frames was 10 seconds long and roughly constitutes about seven gait cycles considering that an average individual takes about 1.4 seconds to complete a gait cycle [10, 20].

the conclusion of Stang [55] cannot be relied upon at least for two reasons. First, the authors studied only five gait templates that were collected from only one individual. Second, Pearson's correlation coefficient between the resultant acceleration of the target and the imitator samples was used to measure the attack's success.

Mjaaland et al. [34–36] also pointed out that the previous studies merely moved beyond minimal-effort mimicry, which led to the investigation of what impact an extensive training would have on the imitation process. The authors hired a total of seven imitators. Six of them were trained for more than an hour. The seventh imitator was trained for as long as 6 weeks. The systematic change in the learning curve of imitators was observed using linear regression. Specifically, $Y(i) = \beta_1 + \beta_2 \times e^{\beta_3/i}$ a non-linear (in β_3) regression model was used, where i represents i th observation, and β_1, β_2 , and β_3 are constants, a non-linear (in β_3) regression model was used [60]. The residual error for the model was defined as $r_i = |Y(i) - y(i)|$, where $y(i)$ is the value at i th observation. The constants β_2 and β_3 and their magnitude indicated the progress and rate of the (un)learning, respectively. The same signs (a downward sloping curve) of these constants implied that the imitators' learning improved. However, opposite signs indicated that the imitators' learning worsened. The authors conducted the experiments under three scenarios: friendly, short-term hostile, and long-term hostile. The *friendly scenario* consisted of collecting regular walking patterns from 50 participants while they walked a fixed distance without any training

or feedback. Each participant was filmed while walking. This scenario was considered a baseline scenario in the study and resulted in an equal error rate of 6.2%. In the *short-term hostile scenario*, one target and six imitators were selected from the 50 participants. The criteria for choosing the six imitators included the stability of their gait (low intra-variation), the **Dynamic Time Warping (DTW)** distance of their gait patterns from the target's gait patterns, and their eagerness to participate in the study. Imitators of both types, whose gait patterns were too close and whose gait patterns were too far from the target, were selected.

Each of the imitators made five attempts, each lasting about an hour. In each session, the video clips of the imitation attempts and distance scores were shown to the imitators. Possible improvements were discussed/suggested before making the next attempt. In the first two sessions, high-level information such as speed, sideways postures, and arm swing was tuned. In the third session, the imitators walked behind the target. This attempt was also recorded and was one of the most valuable (reported by imitators) feedback. Specifically, this attempt helped to synchronize the speed and step length. Session four was focused on finer details such as feet, ankle, heel, toe, hip, torso, and shoulder movements. Session five was the final imitation trial without any training or feedback. Finally, in the *long-term hostile scenario*, one imitator who was available for at least 6 weeks was chosen. The imitator made a total of 60 attempts, which took about 60 hours, 10 hours per week.

The authors concluded that all six imitators could not breach their respective physiological boundaries. The boundaries were visible for each imitator on the graph plotted from the regression analysis. The authors observed that the imitators could adapt specific characteristics of the target but ultimately failed to match all of the traits. The learning curve for some imitators worsened over time. The imitators reported that over-training was making their walk more unnatural and mechanical. The long-term scenario indicated that a person could have multiple plateaus but with a great deal of uncertainty. Various plateaus produced by the same imitator suggested that an imitator can match the target. Still, due to a high level of uncertainty and insufficient data, the authors could not conclude anything.

Interestingly, the authors pointed out that statistical feedback was much more helpful to the imitators than visual feedback. For some attackers, the visual appearance of the walk was almost similar to that of the target. But the data collected from those attempts were far away (DTW) from that of the target patterns. The participating imitators suggested that the most challenging part was to concentrate on the target's different gait factors simultaneously. In addition, they kept forgetting what they had learned in the previous attempt. Ultimately, the authors concluded that gait imitation is challenging to achieve.

Later, Muaaz and Mayrhofer [38] studied WSGait spoofing in two phases: *reenact* and *coincide*. The imitators observed the target user for about 10 minutes while walking next to them in the reenact phase. In the *coincide* phase, real-time feedback was computed by comparing the gait patterns of the target and the imitator. Nine individuals, specifically trained in mimicking body motions and body language, participated in the process. Five of the nine individuals acted as imitators and the rest as victims. None of the imitators could produce gait samples that would match with those of the targeted individuals. The authors finally concluded that circumventing WSGait via spoofing is a difficult task [38].

The experimental setups of the preceding studies imply two underlying hypotheses. First, two visually similar walks from different individuals would produce matching sensor readings or feature values extracted from the sensor readings. Second, human imitators would learn, adapt, and repeat the targeted gait patterns at will. In contrast, Kumar et al. [27] focused on matching features extracted from the sensor readings produced by the target. They further stated that human imitators might be able to learn, adapt, and repeat the targeted gait patterns more easily with the assistance of a device like a treadmill. The attack based on these ideas achieved an average **Spoof False Accept Rate (SFAR)** of 43.66% for accelerometer-based gait authentication over a dataset of 18 users and 1 imitator [27]. Hence, the motivation to explore the idea of Kumar et al. [27] further under improved scenarios.

Besides the works mentioned earlier, some recent studies attempted circumvention of WSGait by intercepting the authentication pipeline and injecting synthetic data [21, 64]. For example, Jia et al. [21] utilized Silhouette

Fig. 1. The pipeline used for training the authentication system of each user. *Genuine* represents those users for whom the authentication system is being trained. *Impostors* applies to every user in the database except the Genuine user. It is important to note that the impostors do not make any deliberate attempt to mimic the Genuine user; their regular gait patterns were used as impostor samples following the past studies [23, 26, 27]. However, *Imitator* is the user who made a deliberate attempt to copy the Genuine user after receiving feedback-based training (see Section 3.3 for more details).

Guided Generative Adversarial Networks (GANs) to spoof video-based gait recognition systems. Zhu et al. [64] used *k*-means clustering to generate a variety of attack patterns and injected them into the pipeline to successfully fool the sensor-based gait authentication system. Our study differs from these studies, as we focus on circumvention via imitation, whereas they investigate circumvention via data injection [63, 64]. Please refer to Section 4.2.2 for more details.

3 DESIGN OF EXPERIMENTS

3.1 Typical Implementation Frameworks for WSGait

A typical implementation of WSGait includes data acquisition via sensors, preprocessing and segmentation, feature extraction and selection, and training and authentication (see Figure 1) [54].

3.1.1 Data Acquisition via Sensors. A variety of wearable sensors, including the accelerometer, gyroscope, and magnetometer available for smart devices, can be used to capture individuals' gait patterns. The accelerometer is considered one of the best sensors to capture the uniqueness of the individual's gait patterns [32]. In addition, the combination of multiple sensors offers better authentication accuracy [31].

3.1.2 Preprocessing and Segmentation. Mechanisms such as moving average and multi-level wavelet decomposition and reconstruction are used to remove noise and outliers from sensor readings [35]. Considering *low computational complexity* and *low recognition latency* as some of the important design goals, the sensor readings are divided into small segments, either gait cycles or frames with a fixed length. Commonly used approaches for detecting gait cycles include local extrema analysis [16–19], zero-crossing detection, and phase analysis, among others. The detected cycles are further aligned and normalized for point-wise comparison. However, frames with fixed length are obtained either in an overlapping or non-overlapping manner. Overlapping frames can be obtained by using a sliding window mechanism [28]. As per previous studies [28, 31], frame length between 8 and 12 seconds and overlap of half of the frame length offer better recognition. A list of features is computed from each frame, and the resultant set of frames is fed to machine learning algorithms for training authentication model. Many state-of-the-art implementation approaches of WSGait rely on frame-based segmentation coupled with machine learning algorithms [11, 30–32, 54].

3.1.3 Feature Extraction and Selection. The existence of this component depends upon the type of matching strategy one follows. For example, no feature extraction or selection is required in case the matching involves a

direct comparison of extracted segments using distance measures such as DTW [12, 37–39] and distance metrics such as Euclidean [16, 18, 19]. The direct comparison usually applies to cycle-based segmentation and is referred to as a template-based approach. However, several recent studies use a fixed-length frame approach. The authors extracted a series of features from the raw sensor readings and employ a variety of machine learning algorithms to train authentication models [11, 23, 31, 32, 41, 59]. In some cases, such as Kumar et al. [27], the authors select the most distinguishing features using **Mutual Information (MI)** between each feature and the class label.

3.1.4 Training and Authentication. The training component primarily involves the training of the authentication model using classification algorithms or the creation of a template for each user from the data collected during the training session. The *authentication* component uses the trained authentication model or the template to classify the incoming test segments (or feature vectors) into genuine or impostor. For training two-class classification algorithms, one needs samples belonging to both genuine and impostor classes. Authors often use samples collected from other users than the genuine user as impostors [23, 27, 31, 32].

3.2 Choice of the Implementation Framework

The fixed-length frame-based approaches [11, 23, 31, 32, 41, 59] coupled with machine learning algorithms (e.g., **k-Nearest Neighbor (kNN)**, **Support Vector Machine (SVM)**, neural network) usually outperform cycle-based approaches that use direct comparison of segmented data points using distance metrics (e.g., Euclidean or Manhattan) or distance measures such as DTW [12, 16, 18, 19, 37–39] in the context of authentication. Thang et al. [59] specifically investigated both approaches and concluded that the former approach achieved significantly better results (92.7%) than the latter (79.1%) on the same dataset. In addition, a recent study by Al-Naffakh et al. [2] presented a qualitative comparison of 29 studies and concluded that the former approach almost always beats the latter. For this work, we implemented both approaches. We found that the latter approach exhibited very high ($\geq 20\%$) error rates on the test data (collected 2 to 3 days later) compared to the former approach that consistently achieved average error rates below 7% (see Zero-effort HTER heatmap in Figure 3). The high error rates exhibited by the latter approach could be an outcome of several factors. The factors include low sampling frequency in the dataset, a relatively more realistic experimental environment (the participant kept the phone in a pant pocket), and the substantially higher number of gait cycles per user (see Table 1). We posit that the authentication system that exhibits more than 20% error rates on the test data would not be suitable for most of the scenarios (e.g., securing expensive cars, houses, offices) in which WSGait is applicable. Needless to say, the models implemented using the latter approach may suffer from even higher error rates if deployed for use in their current form. Therefore, we decided to focus on WSGait implemented using the former approach. Figure 1 depicts the overall experimental setup.

3.2.1 Dataset. The dataset used in this study was originally collected and used by Kumar et al. [27]. Although the dataset consisted of readings from four sensors, Kumar et al. [27] utilized only accelerometer sensor readings, whereas we use the readings of all four sensors. The essential details of the data collection experiment are provided in the following.

The dataset consists of gait patterns collected from 18 Genuine users and 1 Imitator who mimicked each of the Genuine users one by one. The Genuine users walked naturally in a $100 \times 2 \text{ m}^2$ corridor back and forth for about 2 minutes, keeping an HTC-One M8 smartphone in the right pocket of their pants. The phone was always positioned upside-down, and the screen faced the participants' bodies. The particular setup was adopted to keep the data collection environment uniform and consistent among the users. This data collection setup, however, overlooks the scenarios in which the users can place their phone in any position (e.g., upside-up) and location (left or back pocket) on the body. For more details, please refer to Section 4.2.5.

The readings of four sensors, namely the accelerometer, gyroscope, magnetometer, and rotation vector, were captured by an app installed on the phone. The *Genuine* users repeated the exercise after 2 to 3 days. The data

collected in the first phases were used to train and validate the model; thus, it will be referred to as the *Training dataset*. In contrast, the data collected in the repeat session was used for evaluating the model; hence, it will be referred to as the *Testing dataset* in the rest of this article. However, the data collected from the trained *Imitator* (see Section 3.3) will be referred to as the *Mimicry dataset* in the rest of this article.

3.2.2 Preprocessing, Segmentation, Feature Extraction, and Feature Selection. Following previous studies [34–36], a simple moving average technique was applied to filter the noise. Then, the sensor readings were segmented using a sliding window mechanism with a frame length of 10 seconds and an overlap of 5 seconds. The segmentation process resulted in a set of independent frames. On average, there were 20 frames per user in the Training and Testing datasets, whereas there were 66 frames (collected in three independent attempts) for each user in the *Mimicry* dataset. Adapting from previous studies [30, 31, 45, 61], we extracted a series of time and frequency domain features from each of the frames. The specific time-domain features include arithmetic mean; standard deviation; mean absolute change; mean absolute deviation; skewness; kurtosis; mean energy; the number of mean crossings; the number of peaks; first, second, and third quantiles; length of the longest strike below and above the mean; and bin counts in 16 equally thick bins. The bin counts were included primarily because of its effectiveness reported by some of the seminal studies [22, 30], and it provides different information compared to aggregate features. The frequency domain features included the first, second, and third quantiles, and standard deviation of the Fourier transform’s amplitudes.

Each sensor used for data collection has three components: x , y , and z (i.e., three time-series signals). The fourth signal (i.e., magnitude) was computed as $m = \sqrt{x^2 + y^2 + z^2}$ and considered the fourth component. The features were extracted from each component and concatenated together. As a result, we had 136 features from each sensor. Those were too many features to work with, especially if we want to develop a lightweight system. Consequently, following previous studies [14, 29, 53], we included the MI-based feature selector in the pipeline to select the top k most discriminating features. We did not want to have too few or too many features; we experimented with seven different values (start = 20, end = 50, step = 5) of k and decided to use 30 in the end, as it provided the best results under the Zero-effort testing environment. The feature selector computed MI between each feature and the class label and output the top k features with the highest MIs. We observed that the chosen features varied for different users, but not significantly.

3.2.3 Choice of Classification Algorithms. The next step was to choose classifiers to train the authentication models. Following previous studies [23, 26, 30, 31, 40, 45, 57, 59], especially Kumar et al. [27], we employed the **Bayes Network (Bayes)**, **Logistic Regression (LogReg)**, **Multilayer Perceptron (MulPer)**, **Random Forest (RanFor)**, and SVM to classify the feature vectors between genuine and impostor. Besides these classifiers, we included kNN because previous studies [26, 41, 57] have demonstrated its superiority over the other approaches. To be trained, these classifiers required feature vectors from both genuine and impostor classes. The Genuine feature vectors were created from genuine user’s data, whereas the impostor feature vectors were created from data belonging to impostors (see Figure 1). Since there were 17 impostors, we used about 30% (six) of the feature vectors from each, which resulted in a total of 108 impostor feature vectors. Since we had only 20 genuine feature vectors for each genuine user on average, this situation turned into a class imbalance problem. The problem was addressed by oversampling the Genuine feature vectors using SMOTE (Synthetic Minority Oversampling) [8], a widely used method in the field. The SMOTE application helped balance the number of feature vectors in each class and increased the number of feature vectors used for training to 216. As a result, the training matrix for each user model was of size 216×30 for each user authentication model, where 30 is the number of features.

3.2.4 Hyper-parameter Tuning and Performance Evaluation. The hyper-parameter was tuned by conducting 10-fold cross validations on the *Training* dataset using **Half Total Error Rate (HTER)** as a loss function [6]. HTER is defined as an average of **False Accept Rate (FAR)** and **False Reject Rate (FRR)**. The authentication models were evaluated on the feature vectors extracted from the Testing and Mimicry datasets in two phases.

Fig. 2. The human imitator (on the right) is trying to produce accelerometer readings that would closely match that of the target user (on the left) at the feature level. (n) denotes the step # n. The feedback loop (2-3-4-5-2) ends once each of the five features passes the overlap test.

The first phase focused on Genuine testing and was conducted on the genuine feature vectors. The second phase consisted of two parts: *impostor testing* and *imitator testing*. The *impostor testing* was conducted on the impostor feature vectors, which were created from the users other than the genuine users in the *Testing* dataset. This testing environment is referred to as the *Zero-effort* attack environment. However, the *imitator testing* was conducted on the feature vectors extracted from the data available in Mimicry dataset for each genuine user. This testing environment would be referred to as the *Treadmill-assisted* attack environment.

Each of the feature vectors was assigned either genuine (accept) or impostor (reject) class labels during testing. One can devise a policy on how many consecutive accepts or rejects granting or denying access can be devised for different application scenarios. For example, one can secure a smart car via WSGait. The smart-car doors may open only after the user produces five continuous matching frames. For simplicity, like previous researchers, we report the performance of WSGait in terms of FAR and FRR, which are defined as follows:

$FAR = \text{number of impostor frames that were accepted} / \text{number of total impostor frames}$

$FRR = \text{number of genuine frames that were rejected} / \text{number of total genuine frames.}$

FAR and FRR represent Impostor pass and Genuine fail rates, respectively. Similar to FAR (the impostor pass rate), we will use the SFAR to evaluate the Imitator pass rate. In addition, we provide HTER and **Spoof Half Total Error Rate (SHTER)**. The HTER is nothing but an average of FAR and FRR. Likewise, the SHTER is an average of SFAR and FRR [53].

3.3 The Imitation Process

The imitation process is illustrated in Figure 2. The process began by creating an imitator's profile and was followed by training the imitator to mimic the target's gait patterns. The imitator profile was a two-dimensional correlation matrix. The coefficients were computed between the dominant feature set, which consisted of five elements (viz. abs_x, rng_x, nop_x, rng_y, nop_y) and the **Gait Characteristics Adaptable over a Treadmill (GCAT)**. The abbreviations abs, rng, and nop refer to absolute sum, range, and the number of peaks, respectively.

The GCAT included speed, step length, step width, and thigh lift. The data used for computing the correlation was collected at different configurations of GCAT. The speed's configuration values were between 1.2 and 2.8, with an increment of 0.2 miles per hour. The configuration values for the step length included small, regular, and large. Similarly, the step width configurations included close, regular, and wide, and the thigh-lift included front, regular, and back. While changing the configuration values for one characteristic, the levels of other

characteristics were set to regular. The collected data was stored in the *imitator database* along with the corresponding configuration of GCAT.

The dominant features were selected from a list of features that consisted of *abs_x*, *rng_x*, *nop_x*, *api_x*, *bap_x*, *rng_y*, *nop_y*, *eng_y*, *sef_x*, *bap_y*, *sef_y*, *nop_z*, *bap_z*, *sef_z*, *mean_m*, *api_m*, *rng_m*. The abbreviations are as follows: absolute sum (*abs*), range (*rng*), number of peaks (*nop*), energy (*eng*), spectral edge frequency (*sef*), band power (*bap*), and mean (*mean*). It is important to note that these features are substantially different from those that have been used for implementing the authentication systems. Each of these features was examined one by one and added to the dominant feature set (say $\mathbb{D}$) if it was (1) not strongly correlated with any of the features available in $\mathbb{D}$ and (2) strongly correlated with at least one more feature in $\mathbb{L}$ excluding itself. The correlation was considered strong if the $|\text{correlation coefficient}| \geq 0.5$. For example, *abs_x* would be added to $\mathbb{D}$ as it passes both of the conditions as $\mathbb{D} = \emptyset$. Next, *rng_x* would be added to $\mathbb{D}$ because it is not strongly correlated with any features in $\mathbb{D} = \{\text{abs}_x\}$ and is strongly correlated with *bap_x*, *sef_y*, *mean_m*, and *api_m*. Similarly, the rest of the features available in $\mathbb{L}$ were examined and added to $\mathbb{D}$ once they passed the stated conditions. The imitator profile helped to understand how the changes in a specific GCAT affected the values of the individual dominant features.

Figure 2 illustrates the imitator training process for mimicking a specific target. The details of the steps are as follows. (1) Steal WSGait from the targeted user. In practice, the targeted user can be enticed to install an app on her phone (see Section 4.2.1 for more details). The app will collect the desired sensor readings and send the same to a server accessible to the attackers. (2) Search the imitator database and find WSGait that was the closest (with the least mean absolute error) to the *stolen WSGait*. Take a note of the GCAT configurations corresponding to the *closest WSGait* selected from the imitator database. (3) For each of the five dominant features, conduct the overlap test. A feature passes the overlap test if more than 70% of its values (extracted from the mimicked samples) fall inside the range of the values extracted from the target's stolen samples. The process terminates if all five dominant features pass the overlap test simultaneously. (4) In case the overlap test fails for a particular feature, the mean error between the values obtained from *closest WSGait* and *stolen WSGait* for that feature is computed and used along with the highest correlation coefficient corresponding to the feature in the imitator profile to generate the next feedback. For example, let the feature *rng_y* does not pass the overlap test, the mean error is -3.2 , and the correlation between *rng_y* and speed was 0.9 , then the feedback to the imitator would be to *increase the speed*. (5) With the feedback and GCAT configurations, the imitator produces new patterns. The newly obtained patterns are stored in the imitator database along with the updated GCAT configurations. Steps (2) through (5) are repeated until every dominant feature simultaneously passes the overlap test.

Three sets of spoof samples were collected after exiting from the feedback loop. Each set consisted of about 22 frames (see Section 3.1.2). These frames were then classified using the respective authentication models, and corresponding SFARs were calculated.

4 RESULTS AND DISCUSSION

4.1 Results

Results obtained for Zero- and Treadmill-assisted attack scenarios are described next.

4.1.1 Zero-effort Attack Scenario. The first three heatmaps in Figure 3 summarize the performance of WSGait models obtained under the Zero-effort attack scenario. Overall, kNN was the best classifier with the average ZHTER under 3%, followed by RanFor and LogReg with ZHTER of 7%. The rest of the classifiers were able to achieve under 10% error rates. The obtained error rates either closely match or are better than what has been reported in the past under Zero-effort attack scenarios [11, 12, 16, 18, 19, 23, 31, 32, 37–39, 41, 59]. Among the individual sensors (see the first four columns of the ZHTER heatmap), the accelerometer has achieved the best error rates, which aligns with the findings of Lee and Lee [31, 32]. In addition, we observed that WSGait models implemented using multiple sensors exhibited low error rates in general (compare the first 4 and last 11 columns

Fig. 3. The x-axis indicates the combination of sensors. The y-axis represents classification algorithms. Each cell represents percentage error rates for the corresponding combination (classifier and sensors) and is rounded to the nearest integer. A quick comparison between the first (i.e., ZFAR) and the fourth (i.e., SFAR) heatmaps clearly shows the attack's impact. The average FAR has increased from 4% to 26% (see the bottom right corners of the ZFAR and SFAR heatmaps). a, accelerometer; g, gyroscope; m, magnetometer; r, rotation vector.

of the ZHTER heatmap). Although the magnetometer exhibited very high error rates alone, its combination with the accelerometer achieved the overall best error rate (ZHTER of 4%). The results concur with the findings of Lee and Lee [32].

4.1.2 Treadmill-assisted Attack Scenario. The fourth heatmap presents the SFAR, the FAR obtained under the treadmill-attack scenario. Overall, the SFAR heatmap suggests that WSGait models based on kNN were the most affected (SFAR of 41%) ones regardless of the sensor combinations used, followed by Bayes (SFAR of 28%). The rest of the classifiers achieved SFAR between 21% and 22%. SFAR of more than 20% renders the WSGait undesirable for many application scenarios. Therefore, we conclude that treadmill-assisted gait spoofing is a threat to WSGait.

Table 2. Comparison of the Performance of the Attack on the Design of Accelerometer-based WSGait Presented in This and the Previous Study [27]

Classifier	Zero-effort Attack (in %)						Treadmill-assisted Attack (in %)			
	Prev	Curr	Prev	Curr	Prev	Curr	Prev	Curr	Prev	Curr
	FAR	FAR	FRR	FRR	HTER	HTER	SFAR	SFAR	SHTER	SHTER
Bayes	4.75	8.22	14.65	7.91	9.70	8.07	30.14	27.20	22.40	17.56
LogReg	6.16	2.07	6.46	8.87	6.31	5.47	39.16	19.91	22.81	14.39
MulPer	5.90	3.81	6.20	8.66	6.05	6.23	40.23	22.16	23.22	15.41
RanFor	5.90	4.90	3.40	5.88	4.65	5.39	43.55	24.02	23.48	14.95
SVM	5.60	2.89	9.12	8.19	7.36	5.54	38.43	20.96	23.78	14.57
Average	5.66	4.38	7.97	7.90	6.81	6.14	38.30	22.85	23.13	15.38

In the previous study, the authors used a fixed set of 17 features for implementing all user models. In addition, they used the same set of features to train the imitator. In this work, we use 30 highly discriminative features selected from a pool of 136 time and frequency domain features. In addition, we included the feature selector in the classification pipeline that selected different set of features for each user model making it intuitively more difficult to breach. The FAR, FRR, and HTER in the Zero-effort columns suggest that the alternative design of WSGait presented in this work not only achieves a reduced (from 6.81% to 6.14%) error rate but also makes the WSGait more resilient (current SFAR is 22.85% compared to the previous SFAR of 38.30%) to the treadmill-assisted attack. The SFAR of 22.85% is still very high and makes WSGait undesirable for the majority of its application scenarios (see Section 4.2.1).

In addition, the SFAR heatmap suggests that basing a comparison of different design alternatives of WSGait just on the ZFAR, ZFRR, ZHTER, or any metric computed under the Zero-effort attack could be misleading [41, 57].

4.1.3 Comparison with the Previous Study. We compare error rates reported in the previous study and this study in Table 2. We can see that the design alternative presented in this article is more accurate and robust in both Zero-effort and Treadmill-assisted attack scenarios. This is because of the use of a wide variety of features selected based on MI. The SFAR of 22.85%, however, is problematic for several application scenarios of WSGait.

4.2 Discussion

4.2.1 Practicality and Effectiveness of the Attack. Before we discuss the practicality and effectiveness of the attack, we would like to note that the presented attack is proof of concept and can be improved substantially in the future. For example, one can replace a human-plus-treadmill with an electromechanical device for a better calibration of gait characteristics, to obtain better attack results [15]. The practicality and effectiveness of the attacks, such as the one presented in this work, depends on the application scenario, the amount of information (e.g., information about the target, implementation details of the authentication system) that the attackers have access to [25] besides the quality of resources (e.g., the quality of treadmill or a sophisticated humanoid robot) deployed, and how much effort is put in, overall [42]. We describe these factors in the context of the presented attack as follows.

Application scenarios. WSGait may be deployed to protect smart cars [43], lockers, and so forth, or to grant access to critical and highly restricted areas such as secure government offices, cockpits, or military bases [48] as a secondary authentication mechanism if not primary. The possibility of the deployment of gait authentication systems is high because WSGait is considered one of the most challenging traits to mimic [13, 16, 18, 19, 36, 38, 55]. Items such as smart cars and lockers can be stolen and unlocked at attackers' locations using the presented attack method.

Required information to launch the attack. To launch the attack, the attackers need to have access to the gait samples of the target user. The attackers could find numerous ways to steal the required samples. One possible scenario is described next.

The attackers could design an app that collects inertial sensor (requires zero permission [3, 7]) readings and sends the same to a server administered by them. An invite to install the app can be sent to a list of targeted

victims or a list of random victims offering a fantastic shopping deal. The potential victims would install the app to grab the shopping deal. Once attackers receive the target user's data on the server, they can train one or more imitators. Once the attackers are sure that at least one of the imitators is producing the gait patterns that substantially overlap with that of the target's stolen samples, they could launch the presented attack.

Information and resources that would make the attack easier. The more information the attackers have access to, the less effort it may take to launch the presented attack. If the attackers know or can estimate height, weight, body type, gender, and age [1] of the target user and have access to the implementation details such as features or classification algorithms, the likelihood of the attack success can be improved [16]. Moreover, the digital treadmill (Make: LifeSpan; Model: TR1200) used in the presented attack is a simple treadmill with standard belt size and speed control. We believe that advanced treadmills such as the one from NordicTrack (Model: Commercial X32i), which consists of a longer and broader belt, inclination control, push bar, and sled grips, would offer better support to the imitators. With more functionalities, support, and control, additional gait factors such as feet or ankle up-down (exploiting inclination) can be defined, which would only help the imitator.

4.2.2 Advantage over Other Attack Methods. Ratha et al. [47] have described eight possible attack points in a generic biometrics-based system. All but the first attack point rely on either intercepting the authentication pipeline or overwriting/overriding the stored data. Every attack point, except the first, would be generally guarded by the system-level permissions. In other words, the first attack point relies on the user input; thus, it violates the security principles of "Do not trust user input." Needless to say, the presented attack exploits the first attack point. Unless we find a way (e.g., liveness detection in fingerprint or face authentication) to assess the input's validity, it would be difficult to thwart the presented attack.

Although the presented attack requires access to the target's biometric samples, it need not modify the authentication pipeline. In other words, it needs no special privilege. However, to exploit the rest of the attack points, the attacker would require either a write-permission to modify protected areas of the memory or an interception to a highly encrypted authentication pipeline (likely) or both in addition to access to the target's biometric samples [21, 46, 50, 58, 62, 64]. Researchers have argued that once the attackers get permission to modify the target's device, the race is already over [14, 49].

4.2.3 A Numerical Explanation of Working of the Attack? We conducted an in-depth analysis at the feature level to find a numerical explanation behind the working of the attack. Specifically, we analyzed the feature-wise overlap between the genuine user and impostors (the rest of the users in the database) and then between genuine and imitator using the Bhattacharyya coefficient ($\mathbb{BC}$), which is defined as follows [24] and has been used for the same purpose in the past [5, 52]:

$$\mathbb{BC}(p_1, p_2) = \sum_{i=1}^n \sqrt{p_1(i) \times p_2(i)}, \quad (1)$$

where p_1 and p_2 are two discrete probability distributions over the same domain X , n is the number of partitions, and $p_1(i)$ and $p_2(i)$ are the number of data points of the samples p_1 and p_2 in the i th partition. $\mathbb{BC}$ always lies between 0 and 1 (i.e., $0 \leq \mathbb{BC}(p_1, p_2) \leq 1$). Zero denotes no overlap, whereas 1 denotes complete overlap between the probability density functions of the two samples.

The higher the feature overlap is, the more chances of false acceptance. Therefore, we computed the $\mathbb{BC}$ for each of the top 30 features extracted from the data collected from *target* and *impostors* and from *target* and the *imitator*. The $\mathbb{BC}$ between the distributions of features of *target* and *impostors* was found to be ≤ 0.30 for the majority of the top 30 features, whereas $\mathbb{BC}$ was observed to be mostly ≥ 0.30 for the substantially affected *targets* and the *imitator*. For reference, we provide the feature overlap for one of the most, average, and least affected *targets* and the corresponding mimicked samples (Table 3). The table suggests that the success of the attack can be roughly estimated by looking at $\mathbb{BC}$ for most features. The median of $\mathbb{BC}$ s seems to be a good indicator of the SFAR.

Table 3. Demonstrating the Relationship between the Inter-class Feature-level Overlap Indicated by Bhattacharyya Coefficients $\mathbb{BC}(p_{gen}, p_{imp})$ and the SFARs

SN	User11 (SFAR = 61%)		User18 (SFAR = 29%)		User8 (SFAR = 0%)	
	Feature	$\mathbb{BC}$	Feature	$\mathbb{BC}$	Feature	$\mathbb{BC}$
1	npeaks_x	0.25	fftc_std_dev_y	0	stddev_x	0
2	mad_y	0.36	mad_y	0	meanabschange_x	0
3	stddev_m	0.39	stddev_y	0	mad_x	0
4	skewness_z	0.45	mean_energy_y	0	mean_energy_x	0
5	fquantile_y	0.46	fquantile_z	0.04	fftc_std_dev_x	0
6	mean_energy_m	0.46	bin_counts7_z	0.04	fquantile_x	0
7	stddev_y	0.5	strikeabovemean_x	0.06	meanabschange_y	0
8	fftc_std_dev_y	0.54	tquantile_m	0.11	strikeabovemean_z	0
9	mean_energy_y	0.55	meanabschange_y	0.14	ncmean_y	0
10	bin_counts3_z	0.57	fftc_tquantile_y	0.23	mad_y	0
11	meanabschange_y	0.58	fquantile_y	0.23	strikeabovemean_x	0
12	tquantile_y	0.59	mad_m	0.26	mean_energy_m	0
13	bin_counts13_z	0.59	meanabschange_z	0.28	fftc_squantile_y	0
14	fftc_std_dev_m	0.6	amean_m	0.3	bin_counts5_z	0
15	amean_m	0.63	stddev_m	0.3	stddev_y	0
16	strikebelowmean_y	0.63	mean_energy_m	0.34	mean_energy_y	0
17	fquantile_m	0.64	fftc_std_dev_m	0.38	amean_m	0
18	npeaks_z	0.72	fftc_squantile_y	0.39	tquantile_m	0
19	bin_counts13_y	0.75	npeaks_x	0.4	ncmean_z	0
20	bin_counts12_z	0.76	fquantile_m	0.48	tquantile_y	0
21	bin_counts12_y	0.8	skewness_y	0.51	fftc_tquantile_x	0.04
22	bin_counts11_z	0.82	bin_counts12_y	0.51	tquantile_x	0.06
23	npeaks_y	0.85	ncmean_y	0.52	strikeabovemean_y	0.2
24	bin_counts11_y	0.86	tquantile_y	0.58	npeaks_x	0.21
25	bin_counts15_m	0.86	bin_counts11_y	0.62	strikebelowmean_x	0.25
26	npeaks_m	0.89	npeaks_m	0.66	strikebelowmean_y	0.34
27	bin_counts14_y	0.93	bin_counts13_y	0.67	bin_counts13_y	0.43
28	bin_counts15_y	0.94	npeaks_z	0.83	fftc_squantile_x	0.53
29	bin_counts15_x	0.95	strikebelowmean_z	0.88	skewness_y	0.56
30	bin_counts15_z	0.96	npeaks_y	0.93	bin_counts11_y	0.78
31	MedianBC	0.63	MedianBC	0.32	MedianBC	0
32	MeanBC	0.66	MeanBC	0.36	MeanBC	0.11
33	StdDevBC	0.19	StdDevBC	0.27	StdDevBC	0.21

We can see that the MedianBCs of 0.63, 0.32, and 0.0 correspond to 61%, 29%, and 0% SFARs, respectively. Note that the authentication models of different users have different features ranked in the top 30 (see Section 3.2.2 for more details).

4.2.4 Is One Imitator Enough for the High Effort Attack Scenario? A question arises as to whether a single imitator can mimic any individual's gait patterns. The more in-depth analysis of the user-level errors and demographics of the 13 genuine users suggested that the users in the proximity of the imitator's physical characteristics were affected more in general. However, the least affected users were drastically different, in terms of physical traits, from the imitator. The previous study [27] has reported that while training the imitator for particular users, the training process was terminated for some users in less than 10 iterations as the feature overlap passed

the criteria. However, for some users, the feature overlap did not meet the threshold even after 40 iterations. The significant difference between such a user and imitator pair was the height and weight. Thus, we posit that the imitators with similar physical characteristics, especially height and weight, would likely be more successful.

4.2.5 Limitations. Although exhaustive, the experiments presented in this work focused on a widely studied authentication system design that uses frame-based feature extraction and machine learning algorithms. Like most previous studies [36, 38, 55], this study is also limited in terms of the dataset's size. Although the dataset used in this study was collected in a more realistic way than several previous studies, the phone positioning and location were still restricted. In a more realistic scenario, the participant would be allowed to place the phone in the position and location desired [56]. It would be interesting to explore the effectiveness of the attack in that scenario in the future. Furthermore, we would like to note that the data collection process for such investigations is an extraneous exercise for both the researchers and participants. It would be interesting to investigate a less mechanical (e.g., algorithmic) way to study and evaluate the robustness of WSGait [4]. Nevertheless, this work demonstrated that WSGait could be easy to circumvent and draws the research community's attention toward the problem with WSGait.

4.2.6 Possible Countermeasures. Possible countermeasures can include the fusion of sensor reading collected from different devices (e.g., smartphone, smart watch, and smart ring). Researchers may also explore methods that would draw a more robust boundary between the genuine and impostor samples.

5 CONCLUSION

The alarming success (average FAR from 4% to 26%) of the proof-of-concept attack on authentication systems built upon a variety of wearable sensors calls for further research and development of robust countermeasures and rigorous performance testing before WSGait gets deployed for public use. In the future, we plan to explore a dictionary-based attack and possible countermeasures.

ACKNOWLEDGMENTS

We are deeply grateful to the anonymous reviewers for taking the time to review the manuscript and give insightful feedback and comments.

REFERENCES

- [1] Md Atiqur Rahman Ahad, Thanh Trung Ngo, Anindya Das Antar, Masud Ahmed, Tahera Hossain, Daigo Muramatsu, Yasushi Makihara, Sozo Inoue, and Yasushi Yagi. 2020. Wearable sensor-based gait analysis for age and gender estimation. *Sensors (Basel)* 20, 8 (2020), 2424.
- [2] Neamah Al-Naffakh, Nathan Clarke, and Fudong Li. 2018. Continuous user authentication using smartwatch motion sensor data. In *Trust Management XII*, Nurit Gal-Oz and Peter R. Lewis (Eds.). Springer International.
- [3] Zhongjie Ba, Tianhang Zheng, Xinyu Zhang, Zhan Qin, Baochun Li, Xue Liu, and Kui Ren. 2020. Learning-based practical smartphone eavesdropping with built-in accelerometer. In *Proceedings of the Network and Distributed System Security Symposium (NDSS'20)*.
- [4] Lucas Ballard, Daniel Lopresti, and Fabian Monrose. 2007. Forgery quality and its implications for behavioral biometric security. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 37, 5 (2007), 1107–1118.
- [5] Amith K. Belman and Vir V. Phoha. 2020. Discriminative power of typing features on desktops, tablets, and phones for user identification. *ACM Transactions on Privacy and Security* 23, 1 (2020), Article 4.
- [6] Samy Bengio, Christine Marcel, Sebastien Marcel, and Johnny Mariethoz. 2002. Confidence measures for multimodal identity verification. *Information Fusion* 3, 4 (2002), 267–276.
- [7] David Berend, Shivam Bhasin, and Bernhard Jungk. 2018. There goes your PIN: Exploiting smartphone sensor fusion under single and cross user setting. In *Proceedings of the 13th International Conference on Availability, Reliability, and Security (ARES'18)*. ACM, New York, NY.
- [8] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. 2002. SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research* 16 (2002), 321–357.
- [9] Patrick Connor and Arun Ross. 2018. Biometric recognition by gait: A survey of modalities and features. *Computer Vision and Image Understanding* 167 (2018), 1–27.

- [10] Robert Barry Dale. 2012. Clinical Gait Assessment. Retrieved March 24, 2021 from <https://www.sciencedirect.com/science/article/pii/B9781437724110000216>.
- [11] Robertas Damaševičius, Rytis Maskeliūnas, Algimantas Venčkauskas, and Marcin Woźniak. 2016. Smartphone user identity verification using gait characteristics. *Symmetry* 8, 10 (2016), 100.
- [12] Mohammad Omar Derawi, Claudia Nickel, Patrick Bours, and Christoph Busch. 2010. Unobtrusive user-authentication on mobile phones using biometric gait recognition. In *Proceedings of the 2010 6th International Conference on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP'10)*.
- [13] Simon Eberz, Giulio Lovisotto, Andrea Patanè, Marta Kwiatkowska, Vincent Lenders, and Ivan Martinovic. 2018. When your fitness tracker betrays you: Quantifying the predictability of biometric features across contexts. In *Proceedings of the IEEE Symposium on Security and Privacy (SP'18)*.
- [14] Mario Frank, Ralf Biedert, Eugene Ma, Ivan Martinovic, and Dawn Song. 2013. Touchalytics: On the applicability of touchscreen input as a behavioral biometric for continuous authentication. *IEEE Transactions on Information Forensics and Security* 8, 1 (2013), 136–148.
- [15] Lex Fridman. 2020. Walking Robot. Retrieved March 24, 2021 from https://www.linkedin.com/posts/lexfridman_early-version-of-boston-dynamics-humanoid-activity-6698425407957458944-Marp.
- [16] Davrondzhon Gafurov. 2007. Security analysis of impostor attempts with respect to gender in gait biometrics. In *Proceedings of the 2007 1st IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS'07)*.
- [17] Davrondzhon Gafurov and Einar Snekkenes. 2009. Gait recognition using wearable motion recording sensors. *EURASIP Journal on Advances in Signal Processing* 2009 (2009), Article 7, 16 pages.
- [18] Davrondzhon Gafurov, Einar Snekkenes, and Patrick Bours. 2007. Spoof attacks on gait authentication system. *IEEE Transactions on Information Forensics and Security* 2, 3 (2007), 491–502.
- [19] Davrondzhon Gafurov, Einar Snekkenes, and Tor Erik Buvarp. 2006. Robustness of biometric gait authentication against impersonation attack. In *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops*. Lecture Notes in Computer Science, Vol. 4277. Springer, 479–488.
- [20] Heikki Uustal. 2019. Normal Gait. Retrieved March 24, 2021 from <https://medicine.missouri.edu/sites/default/files/orthotics%20review%20case/Normal%20Gait.pdf>.
- [21] Meijuan Jia, Hongyu Yang, Di Huang, and Yunhong Wang. 2019. Attacking gait recognition systems via silhouette guided GANs. In *Proceedings of the 27th ACM International Conference on Multimedia*. 638–646.
- [22] Andrew H. Johnston and Gary M. Weiss. 2015. Smartwatch-based biometric gait recognition. In *Proceedings of the 2015 IEEE 7th International Conference on Biometrics Theory, Applications, and Systems (BTAS'15)*.
- [23] Felix Juefei-Xu, Chandrasekhar Bhagavatula, Aaron Jaech, Unni Prasad, and Marios Savvides. 2012. Gait-ID on the move: Pace independent human identification using cell phone accelerometer dynamics. In *Proceedings of the 2012 IEEE 5th International Conference on Biometrics: Theory, Applications, and Systems (BTAS'12)*. IEEE, Los Alamitos, CA.
- [24] Thomas Kailath. 1967. The divergence and Bhattacharyya distance measures in signal selection. *IEEE Transactions on Communication Technology* 15 (1967), 52–60.
- [25] Hilmi Günes Kayacik, Mike Just, Lynne Baillie, David Aspinall, and Nicholas Micallef. 2014. Data driven authentication: On the effectiveness of user behaviour modelling with mobile device sensors. arXiv:1410.7743
- [26] Rajesh Kumar, Partha Pratim Kundu, and Vir V. Phoha. 2018. Continuous authentication using one-class classifiers and their fusion. In *Proceedings of the 2018 IEEE 4th International Conference on Identity, Security, and Behavior Analysis (ISBA'18)*.
- [27] Rajesh Kumar, Vir V. Phoha, and Anshumali Jain. 2015. Treadmill attack on gait-based authentication systems. In *Proceedings of the IEEE 7th International Conference on Biometrics: Theory, Applications, and Systems (BTAS'15)*.
- [28] Rajesh Kumar, Vir V. Phoha, and Rahul Raina. 2016. Authenticating users through their arm movement patterns. arXiv:1603.02211
- [29] Rajesh Kumar, Vir V. Phoha, and Abdul Serwadda. 2016. Continuous authentication of smartphone users by fusing typing, swiping, and phone movement patterns. In *Proceedings of the IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS'16)*.
- [30] Jennifer R. Kwapisz, Gary M. Weiss, and Samuel A. Moore. 2010. Cell phone-based biometric identification. In *Proceedings of the IEEE 7th International Conference on Biometrics: Theory, Applications, and Systems (BTAS'10)*.
- [31] Wei-Han Lee and Ruby B. Lee. 2017. Implicit smartphone user authentication with sensors and contextual machine learning. In *Proceedings of the IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'17)*.
- [32] Wei-Han Lee and Ruby B. Lee. 2015. Multi-sensor authentication to improve smartphone security. In *Proceedings of the 2015 International Conference on Information Systems Security and Privacy (ICISSP'15)*. 1–11.
- [33] Maria De Marsico and Alessio Mecca. 2019. A survey on gait recognition via wearable sensors. *ACM Computing Surveys* 52, 4 (2019), Article 86.
- [34] Bendik B. Mjaaland. 2009. *Gait Mimicking: Attack Resistance Testing of Gait Authentication Systems*. Master's Thesis. Institutt for Telematikk. <https://ntnuopen.ntnu.no/ntnu-xmloi/handle/11250/261802>.
- [35] Bendik B. Mjaaland. 2010. The Plateau: Imitation attack resistance of gait biometrics. In *Policies and Research in Identity Management*, Elisabeth de Leeuw, Simone Fischer-Hübner, and Lothar Fritsch (Eds.). Springer, Berlin, Germany, 100–112.

- [36] Bendik B. Mjaaland, Patrick Bours, and Danilo Gligoroski. 2011. Walk the walk: Attacking gait biometrics by imitation. In *Information Security*, Mike Burmester, Gene Tsudik, Spyros Magliveras, and Ivana Ilic (Eds.). Springer, Berlin, Germany.
- [37] Muhammad Muaaz and René Mayrhofer. 2013. An analysis of different approaches to gait recognition using cell phone based accelerometers. In *Proceedings of the International Conference on Advances in Mobile Computing and Multimedia*. ACM, New York, NY.
- [38] Muhammad Muaaz and René Mayrhofer. 2017. Smartphone-based gait recognition: From authentication to imitation. *IEEE Transactions on Mobile Computing* 16 (2017), 3209–3221.
- [39] Muhammad Muaaz and Claudia Nickel. 2012. Influence of different walking speeds and surfaces on accelerometer-based biometric gait recognition. In *Proceedings of the 2012 35th International Conference on Telecommunications and Signal Processing (TSP'12)*.
- [40] Claudia Nickel, Holger Brandt, and Christoph Busch. 2011. Classification of acceleration data for biometric gait recognition on mobile devices. In *Proceedings of the Special Interest Group on Biometrics and Electronic Signatures (BIOSIG'11)*.
- [41] Claudia Nickel, Tobias Wirtl, and Christoph Busch. 2012. Authentication of smartphone users based on the way they walk using k-NN algorithm. In *Proceedings of the 8th International Conference on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP'12)*.
- [42] National Institute of Standards and Technology. 2016. Strength of Function for Authenticators - Biometrics (SOFA-B). Retrieved February 8, 2020 from https://www.nist.gov/system/files/documents/2020/07/30/08_newton_biometrics_presentation_final.pdf.
- [43] Kate O'Flaherty. 2020. Apple Watch Security: How an Awesome New Feature Will Help Make CarKey Safe. Retrieved February 8, 2020 from <https://www.forbes.com/sites/kateoflahertyuk/2020/03/30/apple-watch-security-how-an-awesome-new-feature-will-help-make-carkey-safe/#525937ed1ac9>
- [44] Vishal M. Patel, Rama Chellappa, Deepak Chandra, and Brandon Barbello. 2016. Continuous user authentication on mobile devices: Recent progress and remaining challenges. *IEEE Signal Processing Magazine* 33, 4 (2016), 49–61.
- [45] Abena Primo, Vir V. Phoha, Rajesh Kumar, and Abdul Serwadda. 2014. Context-aware active authentication using smartphone accelerometer measurements. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
- [46] Khandaker A. Rahman, Kiran S. Balagani, and Vir V. Phoha. 2013. Snoop-forge-replay attacks on continuous verification with keystrokes. *IEEE Transactions on Information Forensics and Security* 8, 3 (2013), 528–541.
- [47] Nalini K. Ratha, Jonathan H. Connell, and Ruud M. Bolle. 2001. An analysis of minutiae matching strength. In *Proceedings of the International Conference on Audio- and Video-Based Biometric Person Authentication*.
- [48] George I. Seffers. 2018. DISA Moves Beyond Conventional Biometrics. Retrieved February 11, 2019 from <https://www.afcea.org/content/disa-moves-beyond-conventional-biometrics>
- [49] Abdul Serwadda and Vir V. Phoha. 2013. When kids' toys breach mobile phone security. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security*. 599–610.
- [50] Abdul Serwadda, Vir V. Phoha, Zibo Wang, Rajesh Kumar, and Shukla Diksha. 2016. Towards robotic robbery on the touch-screen. *ACM Transactions on Information and System Security* 18, 4 (2016), Article 14.
- [51] Prakash Shrestha and Nitesh Saxena. 2017. An offensive and defensive exposition of wearable computing. *ACM Computing Surveys* 50, 6 (2017), Article 92.
- [52] Terence Sim and Rajkumar Janakiraman. 2007. Are digraphs good for free-text keystroke dynamics? In *Proceedings of the 2007 IEEE Conference on Computer Vision and Pattern Recognition*.
- [53] Zdenka Sitova, Jaroslav Sedenka, Qing Yang, Ge Peng, Gang Zhou, Paolo Gasti, and Kiran S. Balagani. 2016. HMOG: New behavioral biometric features for continuous authentication of smartphone users. *IEEE Transactions on Information Forensics and Security* 11, 5 (2016), 877–892.
- [54] Sebastijan Sprager and Matjaz B. Juric. 2015. Inertial sensor-based gait recognition: A review. *Sensors (Basel)* 15, 9 (2015), 22089–22127.
- [55] Øyvind Stang. 2007. *Gait Analysis: Is It Easy to Learn to Walk Like Someone Else?* Master's Thesis. Department of Computer Science and Media Technology, Gjovik University College.
- [56] Ravichandran Subramanian and Sudeep Sarkar. 2019. Evaluation of algorithms for orientation invariant inertial gait matching. *IEEE Transactions on Information Forensics and Security* 14, 2 (2019), 304–318.
- [57] Chunxu Tang and Vir V. Phoha. 2016. An empirical evaluation of activities and classifiers for user identification on smartphones. In *Proceedings of the IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS'16)*.
- [58] Chee Meng Tey, Payas Gupta, and Debin Gao. 2013. I can be you: Questioning the use of keystroke dynamics as biometrics. In *Proceedings of the Network and Distributed System Security Symposium (NDSS'13)*.
- [59] Hoang Minh Thang, Vo Quang Viet, Nguyen Dinh Thuc, and Deokjai Choi. 2012. Gait identification using accelerometer on mobile phone. In *Proceedings of the 2012 International Conference on Control, Automation, and Information Sciences (ICCAIS'12)*.
- [60] Ronald E. Walpole and Raymond H. Myers. 2012. *Probability & Statistics for Engineers & Scientists*. Pearson Education.
- [61] Changsheng Wan, Li Wang, and Vir V. Phoha. 2018. A survey on gait recognition. *ACM Computing Surveys* 51, 5 (2018), Article 89.
- [62] Zibo Wang, Abdul Serwadda, Kiran S. Balagani, and Vir V. Phoha. 2012. Transforming animals in a cyber-behavioral biometric menagerie with Frog-Boiling attacks. In *Proceedings of the 2012 IEEE 5th International Conference on Biometrics: Theory, Applications, and Systems (BTAS'12)*.

- [63] Benjamin Zi Hao Zhao, Hassan Jameel Asghar, and Mohamed Ali Kaafar. 2020. On the resilience of biometric authentication systems against random inputs. In *Proceedings of the Network and Distributed System Security Symposium (NDSS'20)*.
- [64] Tiantian Zhu, Lei Fu, Qiang Liu, Zi Lin, Yan Chen, and Tieming Chen. 2021. One cycle attack: Fool sensor-based personal gait authentication with clustering. *IEEE Transactions on Information Forensics and Security* 16 (2021), 553–568.

Received June 2019; revised October 2020; accepted December 2020