



HAL
open science

Improving Resilience of Biometric Based Continuous Authentication with Multiple Accelerometers

Tim Van Hamme, Davy Preuveneers, Wouter Joosen

► To cite this version:

Tim Van Hamme, Davy Preuveneers, Wouter Joosen. Improving Resilience of Biometric Based Continuous Authentication with Multiple Accelerometers. 31th IFIP Annual Conference on Data and Applications Security and Privacy (DBSEC), Jul 2017, Philadelphia, PA, United States. pp.473-485, <10.1007/978-3-319-61176-1_26>. <hal-01684348>

HAL Id: hal-01684348

<https://inria.hal.science/hal-01684348v1>

Submitted on 15 Jan 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Improving resilience of biometric based continuous authentication with multiple accelerometers

Tim Van hamme, Davy Preuveneers, and Wouter Joosen

imec-DistriNet-KU Leuven
Leuven, Belgium

{tim.vanhamme, davy.preuveneers, wouter.joosen}@cs.kuleuven.be

Abstract. Biometrics in multi-factor authentication schemes continuously assess behavior patterns of a subject to recognize and verify his identity. In this work we challenge the practical feasibility and the resilience of accelerometer-based gait analysis as a biometric under sensor displacement conditions. To improve misauthentication resistance, we present and evaluate a solution using multiple accelerometers on 7 positions on the body during different activities and compare the effectiveness with Gradient-Boosted Trees classification. From a security point of view, we investigate the feasibility of zero and non-zero effort attacks on gait analysis as a biometric. Our experimental results with data from 12 individuals show an improvement in terms of EER with about 2% (from 5% down to 3%), with an increased resilience against observation attacks. When trained to defend against such attacks, we observe no decrease in classification performance.

1 Introduction

Multi-factor authentication schemes are adopting behavioral biometrics (or biometrics) [3] to continuously verify in the background the identity of users by leveraging information about the user’s device [21, 22], context or the user’s behavior [5, 11] within that context. These trends are often referred to as *Active Authentication*, also known as *Context-aware* [9], *Continuous* [19] or *Implicit* [20] *Authentication*. The key challenges that these authentication schemes aim to address are (1) the ability to conveniently and reliably authenticate the identity of a user, and (2) to continuously assess the confidence in the user’s identity.

One well-known biometric is gait recognition [10, 17] using accelerometer data to analyze motion patterns. While this technique is hardly new, several challenges from a practical feasibility and security point of view remain: (1) there has been little research that investigates the practical resilience of such schemes against sensor displacement, (2) reported high recognition rates were only achieved in a controlled setup where the test subjects are known to walk, making it difficult to ascertain the accuracy – or even the misauthentication resistance – under other conditions and motion activities, and (3) the feasibility and effectiveness of zero and non-zero effort attacks against gait analysis.

Related work	# Sensors	Body position	EER	RR
Mantjarvi et al., 2005 [13]	1	Hip	7%	88%
Gafurov et al., 2007 [7]	1	Hip (back pocket)	7.3%	86.3%
Annadhorai et al., 2008 [1]	1	Ankle	?	84%
Derawi et al., 2010 [6]	1	Hip	20%	?
Nickel et al., 2011 [16]	1	Hip (phone in pouch)	10.3%	?
Ngo et al., 2014 [15]	5 (1 at a time)	Hip, back	14.3%	?
Lu et al., 2014 [12]	1	Different positions	14%	?

Table 1. Comparing the EER and recognition rate of gait authentication schemes

To the best of our knowledge, we are the first to evaluate the effectiveness of using multiple accelerometers to collectively further improve this type of authentication scheme. Additionally, in this work, we enhance the resilience against the above threats with the following contributions:

- We investigate the effectiveness of accelerometers on 9 different places on the body, and analyze the impact of different human activities on the EER.
- We research whether multiple accelerometers can enhance misauthentication resistance, report on the use of different machine learning algorithms, and discuss which combination of on-body positions is the most effective.
- We evaluate a solution that relies on a common set of features, rather than a unique set for each type of activity, to improve classification robustness under diverse circumstances and motion activities.
- We evaluate our authentication scheme against zero-effort and non-zero effort attacks, and compare the results against single accelerometer schemes.

We evaluate our research on the public REALDISP benchmark dataset¹ that was previously collected to evaluate sensor displacement in activity recognition [2, 4]. Based on a study with data from 12 individuals, our results show a recognition improvement reducing the equal error rate (EER) from 5% down to 3%, with an increased resilience against observation and spoofing attacks.

The remainder of this paper is structured as follows. In section 2, we discuss related work on accelerometer based gait authentication. Section 3 presents our multi-sensor approach. In section 4, we describe the experiments and results. We conclude in section 5 summarizing our main insights and discussing future work.

2 Related work

This section reviews relevant research on gait authentication schemes, summarizing the EER and recognition accuracy results in Table 1.

Mantjarvi et al. [13] investigated the feasibility of using gait signals for identification using correlation, frequency domain and distribution statistics. For 36 subjects wearing the accelerometer on 2 different days, correlation proved to be the best method, obtaining a 7% EER and a 88% recognition rate (RR). Similar work by Gafurov et al. [7] compared absolute distance, correlation, histogram, and higher order moments to evaluate performance of the system both

¹ <https://archive.ics.uci.edu/ml/datasets/REALDISP+Activity+Recognition+Dataset>

in authentication and identification modes. Their analysis on 50 subjects showed that the distance metric had the best performance with an EER of 7.3%, and a recognition rate of 86.3%. Annadhorai et al. [1] identified subjects from gait cycles using k-Nearest Neighbor classification. Features were extracted for each gait cycle from accelerometer (3D), pitch and roll data. A subject was identified with an accuracy of 84%. However, these results were obtained on a relatively small data set, with only 2 walks from 4 different subjects. Derawi et al. [6] tested the feasibility of gait as a biometric by using the accelerometer in a smartphone. During an enrollment phase the average gait cycle is determined. Two gait cycles are compared using Dynamic Time Wrapping (DTW). An EER of 20% was achieved on a dataset containing 51 subjects, with two walks per subject. Contrary to previous works, Nickel et al. [16] did not rely on extracting gait cycles to calculate feature vectors, but used Hidden Markov Models to classify gait patterns of 48 subjects. They reported a False Reject Rate (FRR) of 10.42% at a False Acceptance Rate (FAR) of 10.29% (or an EER of $\approx 10.3\%$). A large scale experiment was conducted by Ngo et al. [15] with 744 subjects between 2 to 78 years old, walking under different ground slope conditions. They verified four different gait based authentication methods. The authors conclude that the maturity of the subject's walking ability and the slope greatly influence the performance of gait based user authentication. Lu et al. [12] describe a gait verification system based on Gaussian Mixture Model - Universal Background Model (GMM-UBM) framework. The design objective was to adapt the gait model for mobile phones such that it can account for different body placements and over time variance in the user's gait pattern. The UBM was trained using data from 47 different subjects, the user gait model was tested for 12 subjects. The reported EER was 14%.

3 Challenges with gait authentication schemes

This section identifies challenges and the gap that we aim to bridge when using accelerometer based gait recognition as a biometric in real life scenarios.

3.1 Different body positions and sensor displacement

Most people own at least one mobile device, with different types of sensors. In the future even more sensors will be attached to our body, in the form of smart watches, activity trackers, smart shoes or even smart clothes. Therefore, there is an opportunity to research what positions on the body are the most characterizing and effective for authentication purposes.

However, most of these devices are not fixed at all times to a certain place on our body. They do have an area where they are normally located, but their exact placement varies from time to time, e.g. changing your smartphone from your left to your right pocket, or wearing pants with entirely different pockets. These subtle sensor displacements in the real world, will have an impact on the classification accuracy, and hence the effectiveness of accelerometer based gait authentication schemes.

3.2 Misauthentication resistance under different motion activities

Walking is not the only predominant activity in human life. We are sitting, running, cycling, climbing stairs, etc. as well. A behavioristic should be able to deal with different types of activities. The related work showed that most techniques (1) assume that people are walking and do not consider other activities; (2) explicitly exploit gait cycles to extract features: their first step is always to discover the gait cycle and extract it from the data sample. While the first assumption is reasonable for completely different activities (Wilson et al. [24] achieved an activity classification accuracy of 95%), this is not valid for the latter. It is useless to extract gait cycles for sitting, and not straightforward to find patterns similar to gait cycles for activities like rowing, going to the gym or cycling. Moreover, the related work seemed to struggle when the walking conditions changed slightly (i.e. changing the walking speed, the type of shoes used, the amount of weight being carried, the type of surface and the inclination of the ground). While it might be possible to classify whether the wearer of the accelerometer is running or walking, and maintaining different models for both cases, it certainly is not practical to repeat this for a range of different speeds.

We therefore investigate the feasibility of a common feature set – rather than special features fitted to every particular activity – and the added value of using multiple accelerometers for behavioristic-based authentication. We will use data where the activity is known beforehand. This is reasonable because of high accuracies achieved for activity recognition in other work [24]. Based on our previous research in the field of activity recognition [18], our hypothesis is that we can obtain even higher accuracies for our use case and setting, because our solution does not rely on a fine-grained distinction between activities, as discussed in section 4.4.

3.3 Security threats and attacker model

To evaluate the effectiveness of the proposed scheme, we consider the impact of two different types of attacks:

- **Zero-effort attack:** the adversary is simply another subject in the database that acts as a casual impostor
- **Non-zero effort attack:** the adversary actively masquerades as someone else by mimicking and spoofing the gait pattern of the claimed identity

In the zero-effort attack, we use the data of the other subjects as negative examples for a given user to get insights into the probability of misauthentication.

A non-zero effort attack would occur when the attacker tries to obtain activity patterns of the subject (i.e. observation and spoofing). The attacker attempts to act like the subject by walking at the same pace or mimicking the characteristic activity, as investigated in [14, 8], or he can try to sneak an accelerometer into the coat of the subject. To make these attacks harder to perform, we combine multiple sensors on different places on the body. This way, we collect more data to learn a subject’s movement patterns, with an opportunity to further decrease the EER.

4 Evaluation

This section reports on the experiments conducted to test the concerns expressed in section 3. We use the public REALDISP benchmark dataset [2, 4] to enable the reproducibility of our research results. It contains 17 subjects, all performing 33 different actions, among which: walking, jogging, running, cycling, rowing, etc. All subjects wore 9 sensors on different positions. They performed the set of exercises twice: once with the sensors adjusted carefully by the makers of the dataset; and once adjusting the sensors themselves. The data collected consists of 3D accelerometer, 3D gyroscope, 3D magnetic field measurements and an estimation of the orientation using quaternions. The sampling rate is 50Hz.

4.1 Activity-agnostic behaviometrics

We do not make any assumptions on a particular motion pattern (e.g. presence of gait cycles) so that our behaviometrics can be used for different types of activities.

The REALDISP dataset contains 33 different activities. For each of them we extracted features using the same approach: the data was split in intervals of 128 samples (which is ≈ 2.5 s). For each interval we calculated some straightforward features, in both the time and frequency domain. Among them: mean, standard deviation, kurtosis, mean average derivation, energy in the signal, average resultant vector. This led to a feature vector of length 224. Only the activities walking, jogging, running and cycling had a meaningful amount of samples (≈ 23) per subject. Only 12 subjects appeared to have walking, running and jogging data, of them only 9 cycled. Each subject had performed all actions twice, once with self sensor placement and once with ideal sensor placement.

With authentication in mind, we trained a model for each subject and each activity. We constructed a set which consists for 50% of samples belonging to the subject and for 50% of samples from other subjects. The samples belonging to the other subjects were sampled equally among the total distribution w.r.t. subjects. For each subject adjacent samples w.r.t. time were taken. This set was split in a training and test set using n-fold cross-validation. This process splits the set in n temporal adjacent chunks, in a stratified manner, thus taking into account to what user the samples belong. A model is trained n times, each time leaving a different chunk out for testing. The others are used for training. The number of false positives (fp), false negatives (fn), true positives (tp) and true negatives (tn) are accumulated over the different iterations. This process is repeated for every subject in the dataset. We compared different classification algorithms by calculating the average EER of all body positions for the walking activity (see Table 2). Support Vector Machines produced bad results due to the small amount of training samples compared to the dimension of the feature space. Ensemble methods like AdaBoosting, Random Forests, Bagging and Gradient-Boosted Trees performed a lot better. Because of the robustness w.r.t. outliers in other machine learning experiments and its ability to handle

SVM with sigmoid kernel	SVM with rbf kernel	kNN	GBT	Bagging	AdaBoost	Random Forest
0.511	0.488	0.198	0.068	0.094	0.047	0.036

Table 2. Comparison of EER with different machine learning classifiers

	Walking	Jogging	Running	Cycling
BACK	0.036	0.010	0.040	0.020
LC	0.034	0.033	0.022	0.010
LLA	0.076	0.060	0.029	0.032
LT	0.033	0.021	0.011	0.012
LUA	0.062	0.045	0.019	0.031
RC	0.026	0.025	0.024	0.009
RLA	0.057	0.082	0.061	0.056
RT	0.016	0.018	0.015	0.011
RUA	0.071	0.064	0.031	0.022

Table 3. EERs of ideally placed sensors

heterogeneous features, we decided upon Gradient-Boosted Trees. We will use this model throughout the following experiments.

4.2 Optimal sensor positions on the body

The REALDISP dataset contains data from sensors placed on different positions. This allowed us to evaluate which are the most relevant ones for authentication. We used the approach described above to train a model for each subject. The FAR and FRR can be tuned by demanding a minimal certainty before accepting a sample as genuine. In a first experiment, we used the data collected during walking under an ideal sensor placement. The results were evaluated using 8-fold cross validation. 9 body positions were considered: the back (BACK), the left (LUA) and right upper arm (RUA), the left (LLA) and right lower arm (RLA), the left (LC) and right calf (RC), the left (LT) and right thigh (RT). First, we note that the results are very promising, with really low EERs: $\approx 8\%$ in the worst case scenario, and reducing further down to $\approx 2\%$. Second, the lower body seems to be more relevant than the upper body.

We repeated the same experiment for the other activities: jogging, running and cycling. The EERs are shown in Table 3. The conclusion that the lower body is more informative than the upper body remains valid for the activities considered. This observation holds in all subsequent experiments. Due to the limited amount of data, we cannot make more fine-grained conclusions.

4.3 Impact of sensor displacements

In real life scenarios, a sensor will never be worn on the exact same position. Therefore we investigated the effect of small sensor displacements.

In a first experiment the model was trained with walking samples where the sensors were administered by a professional, while testing with walking data when the sensors were self placed, and vice versa. The results plummeted, with

Position	Walking	Jogging	Running	Cycling
BACK	0.054	0.035	0.059	0.051
LC	0.050	0.036	0.039	0.023
LLA	0.100	0.077	0.052	0.046
LT	0.050	0.043	0.042	0.021
LUA	0.090	0.049	0.035	0.033
RC	0.042	0.039	0.040	0.030
RLA	0.086	0.067	0.054	0.072
RT	0.051	0.027	0.033	0.028
RUA	0.074	0.059	0.050	0.028

Table 4. EERs of slightly displaced sensors (training on both self-placement and ideal placement data) for each body position and different activities

Position	Walking	Jogging	Running	Cycling
RUA & LUA	0.079	0.052	0.055	0.042
RLA & LLA	0.081	0.066	0.064	0.056
RC & LC	0.044	0.052	0.053	0.030
RT & LT	0.055	0.039	0.048	0.025

Table 5. EERs when training on data from both sides of the body

best case EERs of $\approx 45\%$, worst case up to $\approx 50\%$. This can be explained by the lack of walks under sensor displacement in the training set.

A second experiment uses data from both (ideal and self placement) walks as training data. The best EER, for the RC sensor, is $\approx 5\%$. The worst EER is $\approx 10\%$.

The same experiment was repeated for the other activities as well. The results are shown in Table 4. Our earlier conclusion that the upper body seems to be less suited for authentication than the lower body still holds. On top of that, the lower arm consistently has worse EERs than the upper arm. The increase in amount of training data makes our results more consistent.

Fig. 1 illustrates our conclusions. It shows the results w.r.t. EER for the walking activity of the previous experiment (left bar) and the experiment described in section 4.2 corresponding with Table 3 (right bar). It can clearly be seen that in both cases the upper body is less suited for authentication than the




Fig. 1. EERs for walking. The left bar corresponds to Table 4 and the right to Table 3.

Position	Without Cycling	With Cycling
BACK	0.025	0.025
LC	0.020	0.015
LLA	0.055	0.031
LT	0.028	0.016
LUA	0.045	0.029
RC	0.035	0.017
RLA	0.055	0.037
RT	0.024	0.014
RUA	0.043	0.029

Table 6. EERs training on different activities

	BACK	LC	LLA	LT	LUA	RC	RLA	RT	RUA
BACK	0.054	0.033	0.048	0.032	0.056	0.04	0.066	0.049	0.073
LC	0.037	0.052	0.049	0.061	0.067	0.036	0.043	0.049	0.062
LLA	0.051	0.05	0.1	0.05	0.081	0.054	0.083	0.059	0.055
LT	0.033	0.061	0.058	0.053	0.047	0.04	0.048	0.035	0.061
LUA	0.056	0.067	0.08	0.043	0.089	0.065	0.079	0.047	0.068
RC	0.04	0.037	0.057	0.043	0.074	0.044	0.046	0.047	0.061
RLA	0.066	0.042	0.082	0.047	0.08	0.045	0.086	0.051	0.066
RT	0.049	0.052	0.056	0.033	0.047	0.049	0.054	0.053	0.06
RUA	0.071	0.062	0.055	0.064	0.066	0.063	0.063	0.056	0.072

Table 7. EERs when combining multiple accelerometers

lower body and the back. Furthermore, when data of both ideal and self sensor placement (left bar) is used, the EERs suffer a bit, when compared to using data of only ideal sensor placement.

In a third experiment the training set contained the data from the right part of the body and tests were executed with the corresponding left side of the body. 4 fold cross validation yielded EERs of approximately 45%. The clarification is probably similar to the one in the first experiment: there is not enough training data for this type of brutal sensor displacements.

In a fourth experiment the model was trained using data from both the right and left sensor. The tests were conducted using 4 fold cross validation. The results are shown in Table 5. We conclude that this type of sensor displacement has no additional measurable impact.

4.4 Impact of other motion activities

Earlier we argued that having a model for every activity is infeasible. Even more, engineering optimal features for each activity under different circumstances is impossible. We conduct an experiment where we train our model using different activities. In a first experiment we use walking, jogging and running data, for self and ideal sensor placements. The results, using 4 fold cross validation, are shown in Table 6. Compared to training the model using only one activity, as shown in Table 4, the results have improved. The best EER is $\approx 2\%$, while the worst is $\approx 5.5\%$. We assume that using training data at different speeds improves the EER. This needs to be verified using more fine grained data w.r.t. speed.

In the second experiment we added cycling to the dataset, which does not seem to affect the results significantly (see Table 6). The EERs are lower than

in the first experiment, but cycling gave better results in previous experiments as well. Furthermore, only 9 subjects were available for cycling.

4.5 Resilience against observation attacks

To improve the EER and the resilience against observation and spoofing attacks, we combined the data of two accelerometers. A feature vector of length 448 (2×224) is obtained. The experiment is similar as before, using walking data for both self and ideal sensor placement. The results are shown in Table 7. As expected, combining the same sensor yields no new information. The values on the diagonal of Table 7 are similar to the results shown in Table 4. On top of that, the order in which sensors are combined does not matter, since $EER_{i,j} \approx EER_{j,i}$. The best result is achieved by combining the sensors adjusted on the back and the left thigh, which gives an EER of $\approx 3\%$. This is an improvement of $\approx 2\%$ compared to using both sensors separately (see Table 4). However, if we consider for each body position, the results for left and right sensor together, the combination of a sensor placed on a calf with a sensor on the back yields the best results. Furthermore, for each sensor placement, a combination with the back sensor is among the best scoring. Combining two sensors does not always lead to an improvement in performance, i.e. combining the right upper arm and back sensor leads to an EER of $\approx 7\%$, this is higher than the $\approx 5\%$ EER obtained using the back sensor by itself.

For completeness we investigated what would happen if three sensors were used together. This led to a feature vector of length 672 (3×224). We conclude that adding a third sensor leads to a minor improvement, but definitely not in all circumstances. This is illustrated in Fig. 2. The left bar shows the EER corresponding to each body position (as shown in Table 3). For each body position we add the sensor that leads to the best combined EER. The EERs for two sensors are shown by the middle bar. Then the third sensor, leading to the best EER is added, which is illustrated by the right bar.

An attacker can execute an observation attack by collecting accelerometer data through the HTML5 APIs of a mobile browser. An authentication system relying on only one sensor would now be compromised. When two sensors are used, we need to test the feasibility of misauthentication when the attacker constructs a trace, using the obtained data and his own data for the second sensor. We assume that the attacker knows the location of the second sensor.

To test the above use case we trained the system as we did before; using data from the walking activity, where the sensors are placed on the back and the left calf. Positive training samples are combinations of traces from the subject itself. Negative training samples consist of back and left calf data from other subjects. We test the system with genuine combinations of traces and with constructed combinations of traces by an attacker. The attacker combines an obtained back trace and his own left calf accelerometer data. This leads to bad results, an EER of $\approx 17\%$. At a FRR (false rejection rate) of $\approx 10\%$ a FAR (false acceptance rate) of $\approx 37\%$ is achieved. At the threshold used to achieve the result in Table 4, the FAR is $\approx 43\%$.




Fig. 2. Each bar represents the EER of at least one sensor. The left bar shows the EER when using only the first sensor from the description. The middle bar represents EER of the first two sensors. The right bar, is the EER of all three sensors.

In a last experiment we added some samples of the attack to the training data. This led to an EER of $\approx 4\%$, which is only slightly more than before (0.037). The FAR is $\approx 1\%$ when the FRR is $\approx 10\%$. When the old threshold is used, the FAR is $\approx 4\%$. We conclude that the model has to be trained with the observation attack in mind in order to be resilient against it.

The above results were obtained with a public dataset to guarantee the reproducibility of our research results. However, this means there might be some validity threats to generalize our research results. To address this concern, we are collecting bigger datasets to confirm our findings.

5 Conclusion

In this work, we evaluated the resilience of the accelerometer as a biometric for gait authentication, and more specifically the effect on the equal error rate (EER) when using multiple sensors at different places on the body.

Our experiments on gait authentication with a single accelerometer using Gradient-Boosted Trees and a fairly elaborate feature vector showed low EERs and recognition accuracies that go beyond the state-of-the-art. For data collected from 12 subjects and on 9 different places on the body, we obtained EER values between 2 to 8% for ideally positioned sensors. However, further experiments demonstrated that the accuracy dropped significantly after subtle sensor displacements, with the EER worsening from single digits to about 45%. We obtained similar results when sensors were displaced from one side to the other side of the body. When we incorporate data from displaced sensors in our training data set, the accuracy improves again with EERs between 5 to 10%, i.e. slightly worse compared to our first experiment. We also measured the impact of other motion activities, including jogging, running and cycling, and their effect on misauthentication.

We evaluated the effectiveness of multiple accelerometers for gait authentication and the impact on the classification accuracy, demonstrating further

improvements in the EER of $\approx 2\%$. Additionally, as a hacker can carry out an observation attack and obtain accelerometer data with the HTML5 APIs, we investigated the resilience of our multi-sensor scheme against spoofing attacks. Our experimental results show our scheme is more robust against such attacks under the condition that not all sensors are compromised.

As future work, we will investigate to what extent motion patterns are independent. We will execute additional attack experiments based on an adversary leveraging his own sensor data to reproduce traces of the victim in order to ascertain whether such attacks are feasible and practical. Furthermore, we will investigate whether it is feasible to modularize our multi-sensor behavioristic-based authentication scheme, not by fusing the different data sets before training but rather fusing individual decisions based on each data set [23], allowing for more flexibility to combine different behavioristics at runtime.

Acknowledgments

This research is partially funded by the Research Fund KU Leuven and DiskMan. DiskMan is a project realized in collaboration with imec. Project partners are Sony, IS4U and Televic Conference, with project support from VLAIO (Flanders Innovation and Entrepreneurship).

References

1. Annadhorai, A., Guenterberg, E., Barnes, J., Haraga, K., Jafari, R.: Human identification by gait analysis. In: Proceedings of the 2nd International Workshop on Systems and Networking Support for Health Care and Assisted Living Environments. pp. 11:1–11:3. HealthNet '08, ACM, New York, NY, USA (2008)
2. Baños, O., Damas, M., Pomares, H., Rojas, I., Tóth, M.A., Amft, O.: A benchmark dataset to evaluate sensor displacement in activity recognition. In: Proceedings of the 2012 ACM Conference on Ubiquitous Computing. pp. 1026–1035 (2012)
3. Bailey, K.O., Okolica, J.S., Peterson, G.L.: User identification and authentication using multi-modal behavioral biometrics. *Computers & Security* 43, 77 – 89 (2014)
4. Baños, O., Tóth, M.A., Damas, M., Pomares, H., Rojas, I.: Dealing with the effects of sensor displacement in wearable activity recognition. *Sensors* 14(6), 9995–10023 (2014)
5. Crossler, R., Johnston, A., Lowry, P., Hu, Q., Warkentin, M., Baskerville, R.: Future directions for behavioral information security research. *Computers and Security* 32, 90–101 (2013)
6. Derawi, M.O., Nickel, C., Bours, P., Busch, C.: Unobtrusive user-authentication on mobile phones using biometric gait recognition. In: 2010 Sixth International Conference on Intelligent Information Hiding and Multimedia Signal Processing. pp. 306–311 (2010)
7. Gafurov, D., Snekkenes, E., Bours, P.: Gait authentication and identification using wearable accelerometer sensor. In: 2007 IEEE Workshop on Automatic Identification Advanced Technologies. pp. 220–225 (2007)
8. Gafurov, D., Snekkenes, E., Bours, P.: Spoof attacks on gait authentication system. *IEEE Transactions on Information Forensics and Security* 2(3), 491–502 (2007)

9. Hayashi, E., Das, S., Amini, S., Hong, J., Oakley, I.: Casa: Context-aware scalable authentication. In: Proceedings of the Ninth Symposium on Usable Privacy and Security. pp. 3:1–3:10. SOUPS '13, ACM, New York, NY, USA (2013)
10. Kale, A., Cuntoor, N., Yegnanarayana, B., Rajagopalan, A.N., Chellappa, R.: Gait analysis for human identification. In: Proceedings of the 4th International Conference on Audio- and Video-based Biometric Person Authentication. pp. 706–714. AVBPA'03, Springer-Verlag, Berlin, Heidelberg (2003)
11. Kayacik, H.G., Just, M., Baillie, L., Aspinall, D., Micallet, N.: Data driven authentication: On the effectiveness of user behaviour modelling with mobile device sensors. CoRR abs/1410.7743 (2014)
12. Lu, H., Huang, J., Saha, T., Nachman, L.: Unobtrusive gait verification for mobile phones. In: Proceedings of the 2014 ACM International Symposium on Wearable Computers. pp. 91–98. ISWC '14, ACM, New York, NY, USA (2014)
13. Mantyjarvi, J., Lindholm, M., Vildjiounaite, E., Makela, S.M., Ailisto, H.A.: Identifying users of portable devices from gait pattern with accelerometers. In: Proceedings. (ICASSP '05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005. vol. 2, pp. ii/973–ii/976 Vol. 2 (2005)
14. Mjaaland, B.B.: The Plateau: Imitation Attack Resistance of Gait Biometrics, pp. 100–112. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
15. Ngo, T.T., Makihara, Y., Nagahara, H., Mukaigawa, Y., Yagi, Y.: The largest inertial sensor-based gait database and performance evaluation of gait-based personal authentication. *Pattern Recognition* 47(1), 228 – 237 (2014)
16. Nickel, C., Busch, C., Rangarajan, S., Möbius, M.: Using hidden markov models for accelerometer-based biometric gait recognition. In: 2011 IEEE 7th International Colloquium on Signal Processing and its Applications. pp. 58–63 (2011)
17. Ntantogian, C., Malliaros, S., Xenakis, C.: Gaithashing: A two-factor authentication scheme based on gait features. *Computers & Security* 52, 17 – 32 (2015)
18. Ramakrishnan, A.K., Preuveneers, D., Berbers, Y.: A modular and distributed bayesian framework for activity recognition in dynamic smart environments. In: Ambient Intelligence - 4th International Joint Conference, AmI 2013, Dublin, Ireland, December 3-5, 2013. Proceedings. pp. 293–298 (2013)
19. Shepherd, S.: Continuous authentication by analysis of keyboard typing characteristics. In: European Convention on Security and Detection. pp. 111–114 (1995)
20. Shi, E., Niu, Y., Jakobsson, M., Chow, R.: Implicit authentication through learning user behavior. In: Proceedings of the 13th International Conference on Information Security. pp. 99–113. ISC'10, Springer-Verlag, Berlin, Heidelberg (2011)
21. Spooren, J., Preuveneers, D., Joosen, W.: Mobile device fingerprinting considered harmful for risk-based authentication. In: 8th European Workshop on System Security, EuroSec 2015, France, April 21, 2015. pp. 6:1–6:6 (2015)
22. Spooren, J., Preuveneers, D., Joosen, W.: Leveraging battery usage from mobile devices for active authentication. *Mobile Information Systems* 2017, 14 (2017)
23. Tao, Q., Veldhuis, R.: Threshold-optimized decision-level fusion and its application to biometrics. *Pattern Recogn.* 42(5), 823–836 (2009)
24. Wilson, J., Najjar, N., Hare, J., Gupta, S.: Human activity recognition using lzw-coded probabilistic finite state automata. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 3018–3023 (2015)