Opposing Data Exploitation: Behaviour Biometrics for Privacy-Preserving Authentication in IoT Environments

ANDRAŽ KRAŠOVEC, European Comission, Joint Research Centre (JRC), Italy and University of Ljubljana, Faculty of Computer and Information Science, Slovenia

GIANMARCO BALDINI, European Comission, Joint Research Centre (JRC), Italy

VELJKO PEJOVIĆ, University of Ljubljana, Faculty of Computer and Information Science, Slovenia

Multimodal data harvested by the Internet of Things sensors has recently been utilised for behavioural biometrics and consequently user authentication. While strengthening the security, these data nevertheless present a privacy threat to users whose behaviour can now be modelled in detail, thus allowing the authenticating authority to not only know *who* is present, but also *what* the present person is doing. In this work we reconsider IoT sensor-based authentication and provide a solution mitigating the privacy risk associated with unnecessary information leaks. Our approach harnesses adversarial learning and identifies such a projection of the data that maintains identity separability, yet obfuscates activity separability, thus ensuring that the authenticating authority can successfully identify the user, but not her actions within the sensed environment. We evaluate our approach on a real-world dataset of three activities performed by fifteen users and show that the activity obfuscation is achieved without compromising identification capabilities.

CCS Concepts: • **Computing methodologies** → **Neural networks**; • **Security and privacy** → *Authentication*; • **Human-centered computing** → Ubiquitous and mobile computing.

Additional Key Words and Phrases: Behavioural Authentication, Adversarial Learning, Privacy Preservation

ACM Reference Format:
Andraž Krašovec, Gianmarco Baldini, and Veljko Pejović. 2021. Opposing Data Exploitation: Behaviour Biometrics for Privacy-Preserving Authentication in IoT Environments. In *The 16th International Conference on Availability, Reliability and Security (ARES 2021), August 17–20, 2021, Vienna, Austria.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3465481.3470101

1 INTRODUCTION

For seamless password-less authentication, Alice is using a behaviour-based IoT authentication system, developed by Carl. IoT devices in such a system are equipped with a variety of sensors, which can be used to monitor, record, and transmit Alice's activities and even confirm her identity based on her behavioural patterns. Yet, Carl is interested in many Information Technology (IT) related sectors, one of which is online advertising. He realised that he can use this (presumably authentication purposed) IoT data to reveal Alice's daily routines and profile her behaviour for targeted advertisement. Alice does not want to go back to using other forms of authentication like passwords, yet she is uncomfortable with sharing data with Carl, aside for those necessary for an authentication system to work.

Moving away from traditional authentication tools, behavioural biometrics present an attractive alternative, as they remain more convenient than passwords, yet alleviate privacy concerns associated with "hard" biometrics [8]. Nevertheless, the concern of data abuse still remains, as environment sensing can provide much more information

Manuscript submitted to ACM

than needed for person identification (e.g., it can be used to detect the activities of the user as well). Therefore, in this work we focus on reducing the amount of auxiliary information released by IoT environments, while preserving the identification capabilities of the data exported to the authenticating service. The main contributions of our work include:

- We devise an adversarial learning solution capable of obfuscating information pertaining to a user's activity in an office from IoT-sensed data, to the point of different activities being indistinguishable by an adversary, while retaining the identification accuracy delivered by an authentication mechanism using the said data.
- We demonstrate that privacy retention mechanisms and user identification accuracy are not mutually exclusive. Our adversarial learning solution is able to achieve higher identification accuracy, compared to an identical counterpart without an adversarial element, while simultaneously ensuring increased privacy.
- Exploring an alternative scenario where the goal is to achieve identity-secluding activity recognition, we demonstrate the generalisability of the proposed approach.

In this paper we present our initial exploration of privacy-preserving authentication[1] in IoT. This preliminary work does not represent a full-fledged authentication system, nor have we put the necessary effort to achieve high-accuracy user identification. Instead, we believe that work presented here can provide a crucial link bridging the privacy of a user's local environment and the capabilities offered by remote authentication services. Therefore, with the help of the technique presented in this paper, we envision the potential for the adoption of behavioural biometric authentication in standard authentication protocols such as *OAuth*[2].

2 RELATED WORK

Widespread adoption of sensors such as fingerprint scanners and cameras rendered biometric authentication a viable form of authentication [18]. However, biometric data tend to be sensitive, thus raising privacy issues associated with data collection and storage [23].

Behavioural biometric authentication built upon non-sensitive sensor data delivers a possible solution to the privacy concerns raised by "hard" biometrics [12]. Touch and keyboard dynamics, for instance, are commonly researched behavioural traits [16], as it turns out that individual differences among users can be observed from touch screen and keyboard [5, 17] interaction patterns. Additional modalities, such as acceleration or gyration, can also be sensed and incorporated in such systems to improve classification performance [19]. Gait dynamics, where a user's identity is inferred based on his or her walking pattern is another popular behavioural biometric approach [7]. Even "softer" sensors, such as PC or smartphone usage information, can be used for behaviour-based authentication [10, 13]. Finally, since the opportunities for authentication remain limited to moments when a user is actively expressing behaviour captured by the corresponding sensor – e.g. to moments when a user is walking, if a gait sensor is used – multimodal behaviour biometric systems fusing data from different sensors have been proposed [21].

Behavioural biometrics bring us a step closer towards privacy-sensitive authentication, yet current implementations fail to acknowledge that seemingly benign sensor data can be misused for malicious purposes. For instance, while data from an accelerometer placed near a keyboard can be used for authentication (i.e. by deducing a user-specific key press

[1]The terms user identification and authentication have a different meaning in literature. User identification is the process where the recognition system determines the identity of a user by comparing the behavioural biometrics of the user with the set of the behavioural biometrics of all the users in a data set. User authentication is the process of confirming the claimed identity of a user against another user using the behavioural biometrics data and it can be used to mitigate a masquerading attack when a malicious user attempts to masquerade as a legitimate user. For the sake of simplicity in this paper we use the term identification and authentication with the same meaning, which relates to the user identification process. In addition, we highlight that the study presented in this paper is based on a closed data set and supervised classification is used.
[2]https://www.oauth.net/

intensities and tempo), the same data can also be used to reveal what the person is typing [14]. Similarly, gait-based authentication harnesses data that can be used for fine-grain activity detection.

Recently, adversarial learning emerged as a means of ensuring "domain filtration" in sensor data. The technique incorporates an adversary in a neural network, either as a classifier or a regressor or as a deceitful input generator. After the initial discovery in 2014 by Goodfellow et al. [6], generative adversarial networks (GAN) became primarily popular as they provide the ability to successfully generate artificial media, such as images, music, or video [1, 4]. Our work, on the other hand, is inspired by Liu et al.'s [11] privacy adversarial network (PAN). Instead of providing adversarial input like GANs, PAN presents an adversarial classifier that is used to infer information beyond the initial intent of the application. Such adversarial classifier provides feedback to the encoder input of the network to better filter out unwanted information before passing it to other parts of the network. The authors demonstrate that image classification is possible even with the obfuscation of original images. The approach is not only able to render the images unrecognisable by a human, but is also improving the classification accuracy compared to a non-adversarial learning approach. These findings coincide with the results presented in our paper.

3 PRIVACY ENHANCING ADVERSARIAL LEARNING

Data purposefully gathered for user identification may include further information about users' current activity, e.g. whether they are reading a document on a PC, browsing the Internet, typing on the keyboard, or even performing actions around the environment. To mitigate such data leaks, we harness adversarial learning techniques to filter out information unrelated to user identification. Therefore, our system incorporates an Artificial Intelligence (AI) attacker in a form of a neural network that is integrated into the training process. Its goal is to exploit gathered IoT data in order to recognise activities in an office setting and consequently give feedback to the encoder, which in turn then adapts the data it passes through, so that the amount of information pertaining to user activities is reduced.

3.1 Threat Model

In this work we strive to protect users from multi-purpose data inference of authorities that, besides providing legitimate biometric behavioural authentication services, infer other information, e.g. user's *current* activity. Therefore, we focus on minimising the accuracy of the activity classifier while retaining the user identification capabilities of a system. The user is still required to provide the initial raw training data and we assume that the training process can take place at any location, including the authority's servers. However, we assume that such data is not sensitive as it relates only to *past* activities. Once trained, the encoder part of our system is placed between sensor devices and the Internet gateway, ensuring that only identification-related information is sent outside the local network of IoT devices. We assume that the adversary does not have access to the local network and the data transmitted in such a network.

3.2 Privacy Adversarial Learning Network Architecture

To implement a system which is capable of identifying users while obfuscating information, which can be exploited for a privacy attack, we design a three-part neural network (displayed in Figure 1). It consists of an entry point encoder network, and user and activity classifiers, predicting the user's identity and the current activity respectively. In essence, the concept of the complete system is to train the encoder to produce features solely connected to user identification, while obfuscating other informative aspects of the data, which can be used to discriminate the current user activity.

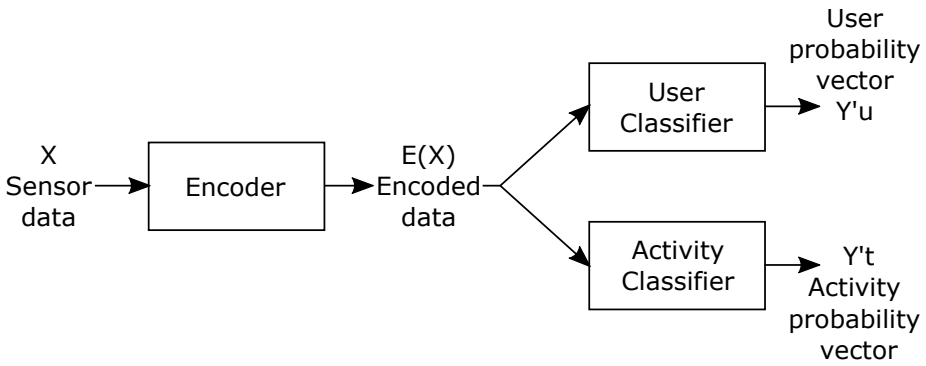

Fig. 1. Privacy-preserving adversarial network for authentication in IoT environments.

3.2.1 Encoder. A multi-layer encoder network serves as the entry point of our system. As we utilise time sequential data, the encoder consists of a Long Short-Term Memory (LSTM) layer, and three subsequent fully connected layers. Its purpose is to take the $\mathcal{X}$ input data array and perform a dimensionality reduction transformation into encoded data $E(\mathcal{X})$. This lower dimensionality feature set retains only user identification relevant information and serves as the input of both classifiers.

3.2.2 User Classifier. User classifier consists of multiple fully-connected layers. It takes $E(\mathcal{X})$ encoded data as an input to output the probability vector of inferred user classes Y'_u.

3.2.3 Activity Classifier. Activity classifier on the other hand tries to extract activity information from the available encoded data $E(\mathcal{X})$. Similarly to the user classifier, it is a dense deep neural network, which provides a probability vector of activity predictions Y'_t.

3.3 Privacy Adversarial Learning Training Algorithm

We guide the training algorithm in a way that only information that is relevant for user identification is processed through the encoder. Starting from the algorithm presented in [11], we split each training epoch into three separate parts, each concerned with training a specific aspect of the network, as presented in Figure 2. We utilise the categorical cross entropy loss as a performance metric in each of the three training parts.

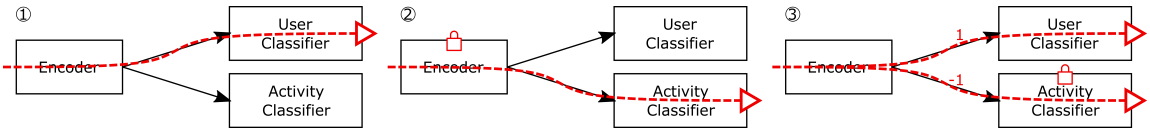

Fig. 2. Training steps of the privacy-preserving adversarial learning network; 1) training of the encoder and the user classifier, 2) training of the activity classifier through the encoder with locked parameters, and 3) training of the complete system, with the negative loss function weight and locked parameters of the activity classifier.

(1) **User classifier training**; in every training epoch, we first train the encoder and the user classifier. The goal of this step is to generate features relevant for the classification of users and ensure that the user inference classifier predicts user classes as accurately as possible. The user class labels are one-hot encoded, thus, we can discard

predicted probability of false user classes. Hence, the simplified categorical cross entropy loss function over N samples with the y'_u probability of the true predicted class:

$$\frac{1}{N}\sum_{n=1}^{N} -\log y'_{u_n} \tag{1}$$

(2) **Activity classifier training**; the goal of the adversarial training step is to train the adversarial part of the network to predict the user's current activity, thus relaying knowledge on how well the encoder obfuscates activity relevant information. We do not want to adapt the encoder to include activity relevant information in the generated features, therefore we freeze the weights of the encoder for the duration of this training step. Similarly to the previous loss function, we simplify the categorical cross entropy loss function, replacing the user with the probability of the true activity class y'_t:

$$\frac{1}{N}\sum_{n=1}^{N} -\log y'_{t_n} \tag{2}$$

(3) **Combined loss function training**; in the last step we train the complete network. To properly instruct the training algorithm to retain only relevant information while obfuscating the rest, we weight the loss function of user inference and activity inference classifiers with 1 and -1 respectively. Additionally, we freeze the weights of the activity classifier, otherwise we would artificially hurt its performance with a negatively weighted loss. The resulting loss function is therefore the difference of the previous categorical cross entropies over N samples:

$$\frac{1}{N}\sum_{n=1}^{N} -\left(\log y'_{u_n} - \log y'_{t_n}\right) \tag{3}$$

4 IMPLEMENTATION AND EXPERIMENTAL EVALUATION

4.1 Dataset

We use a publicly available dataset collected in our previous research [9] that consists of sensor data gathered by IoT devices placed in an office-like environment. The dataset was collected with 21 users (we retain 15) performing three different long-lasting activities. To eliminate confusion, we define some terms related to our dataset. *Activity* is referring to the three different scenarios users have to complete during the user study. Every activity is performed twice by every user, resulting in two activity *runs* per user per activity. Regardless of the activity, every run is specified as a *session*.

Each activity the participant has to perform, serves a different purpose and is designed to excite a distinct set of sensors. The first activity requires a user to type and send an email using a PC terminal placed in the office. The second activity is less linear and requires the user to check weather forecast for a locality of her choice and recommend tourist attractions on the basis of the forecast. The third activity is devised as a treasure hunt and the participant needs to navigate around the environment, follow instructions placed in a series of boxes, and conclude the session by generating a graph on a PC using data found in the last box.

The collected dataset consists of sixteen hours of IoT data (15M records) gathered from different sensors, placed around the environment, namely: (a) accelerometer and gyroscope, (b) force sensors, (c) Passive InfraRed (PIR) sensors, (d) Hall sensors, (e) PC resources monitor. Each sensor is strategically placed to capture various users' behavioural patterns to accomplish identification as accurately as possible. The sample rate of the accelerometer, gyroscope, and force sensors is 200Hz, PC resources monitor gather data once every second, while the PIR and Hall sensors only report on state change (e.g., presence/absence of magnetic field).

4.2 Data Filtering and Processing

Due to suspiciously short session duration, or in other cases missing data from certain sensors, presumably related to not following instructions and issues in data capture process respectively, we filter out the data of six users. Therefore, we retain data of the remaining fifteen participants, who each complete all six sessions successfully.

To assess whether our identification system generalises well over multiple sessions, we take the best practice approach trained on a selected set of sessions and then test it with previously unseen sessions [3]. More specifically, we utilise all first runs of each activity to train the system, then use the remaining, previously unseen activities to evaluate the performance of the generated model. With the data split determined, next we segment the data into one second time intervals and separately for training and testing set generate both time and frequency domain features, such as mean, median, standard deviation, mean crossing rate, and mean derivative value in the time domain, and power spectrum of the frequency domain, obtained by applying discrete Fourier transform with a Hamming window. In the end we retain about 41k data rows, 24k of which are for training purposes.

Last processing step we perform is the hyperparameter tuning of the system. In succession we utilise a random and grid search of the predetermined hyperparameter space, including learning and dropout rate, number of fully connected layers in the encoder, size of the LSTM layer, and number of previous timesteps we take into account.

4.3 Baseline Approaches

We compare our privacy-preserving adversarial learning system with both traditional and deep learning-based user identification baseline approaches.

4.3.1 Traditional Machine Learning Algorithms. IoT environments most often consist of computationally limited low-power devices. Therefore, we select lightweight traditional machine learning algorithms that could be implemented on a low-powered IoT device. We compare our approach against K-Nearest Neighbours (KNN), Linear Discriminant Analysis (LDA), and Naive Bayesian (NB) classifiers. Moreover, some traditional algorithms that are proven to perform well on user classification problems and could run on a more powerful server device (similarly to our adversarial network) are ensemble algorithms, from which we choose Random Forest (RF) and AdaBoost (ADA) classifiers.

4.3.2 Detached Classifiers for User Identities and Activities. To isolate the impact of the privacy-preserving aspect of our adversarial learning approach, we also compare it to a detached version of itself. Specifically, we take our devised architecture, and separately train it for user and activity classification. We expect to see a marginal improvement in user accuracy and a significant improvement in activity accuracy, compared to the joint adversarial learning network model.

5 RESULTS

Privacy-preserving authentication aims to jointly achieve high accuracy for user identification task and low accuracy for activity recognition task, where the latter classifier is indeed trained to achieve the best possible activity recognition accuracy. Therefore, in this section we evaluate *the accuracy of user identification* for our deep learning-based user identity classifier against a range of alternatives, and compare *the loss in accuracy of activity classification* induced by our adversarial approach. Furthermore, we also examine the scalability of our approach as we increase the number of users supported by the system. Finally, we assess the potential of our solution for tackling an inverted problem, where the activity needs to be recognised, while the identity of a user performing the activity needs to remain hidden.

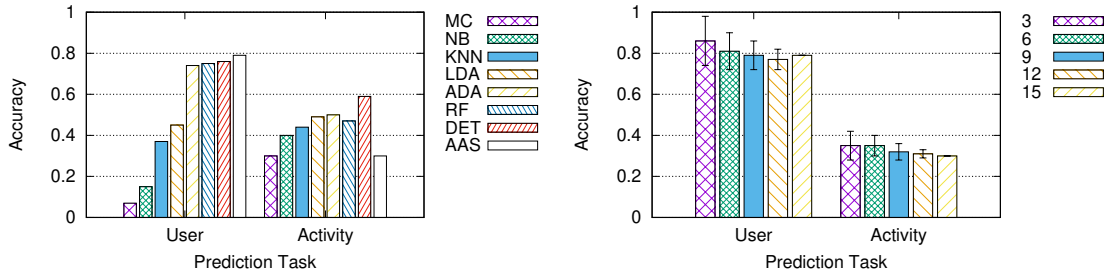

Fig. 3. LEFT: Accuracy comparison of different learning algorithms for user and activity classification. We observe that our detached network (DET) outperforms all shallow approaches, while the adversarial privacy-preserving network (AAS) achieves even higher user classification accuracy. RIGHT: User and activity classification accuracy for different number of users. Error bars represent standard deviation. We observe a decrease in user classification accuracy with the rising number of users, as well as the decrease in variance, due to the lower effect of a single user on the performance. The same effect (with the inverse target result) can be observed for activity inference, as the algorithm more successfully obfuscates activity-related information if data come from a higher number of users.

5.1 Comparison To Baseline Approaches

On the left side of Figure 3, we show the results of user and activity classification accuracy of different classifiers. All models, with the exception of our adversarial learning solution, are separately trained to provide the best possible accuracy for both, user and activity classification, while our privacy-preserving adversarial learning network is trained as described in Section 3. We see that shallow-learning classifiers, including KNN, Naive Bayes and LDA achieve much poorer user classification accuracy than ensemble algorithms and deep learning approaches. Furthermore, our adversarial learning system achieves the highest user classification accuracy, despite being built upon encoded data that obfuscates activity information. In activity classification, traditional models perform poorly all around, achieving classification accuracy lower than 50%. The detached deep learning classifier outperforms other approaches achieving almost 60% accuracy. Yet, our adversarial learning approach successfully obfuscates information pertaining to activity to the point where the accuracy of activity classification drops to that of the majority classifier.

5.2 User Scalability Performance

An authentication system's applicability is tightly related with its ability to adapt to an expected number of potential users. Password-based systems can potentially scale to an infinite number of users. The same does not hold for classifier-based approaches, such as ours, as with the rising number of users the system is presented with more choices to choose from. Therefore, we analyse the drop in user identification accuracy and the ability of the model to obscure activity-related information with the rising number of users.

We sample three, six, nine, twelve, fifteen users and generate a model based on that users' data. We ensure stability of the results by sampling each number of users as many times as it takes for every user to appear in 80 separate models. We plot the average user and activity inference accuracy on the right side of Figure 3.

As expected, the accuracy falls with the rising number of users, from 86% with three users, down to 77% with twelve users. The opposite goes for the system's ability to obfuscate data of activity related information, as we see an increase in accuracy of the adversary activity inference classifier with the lower number of users, as with the lower number of users, there is less chance for inter-user intra-activity reduction in variance. Moreover, we observe the varying user

discernibility, as the standard deviation of user and activity accuracy increases with a lower number of users, as with a lower number of users, each user has a larger impact on the accuracy of the system.

5.3 Inverted Classifiers Learning

Our work is motivated by the need to restrict the authenticating service from uncovering information that is not necessary for the identification and authentication process. Nevertheless, an orthogonal problem is also common in IoT environments. For example, activity recognition is crucial in domains such as elderly care [24], yet one can envision situations where the identity of an elderly user should not be revealed to the service performing activity recognition.

Our modular approach easily enables role inversion of our classifiers. Thus, we now utilise the user identity inference network as an adversary and inspect whether our algorithm is able to distinguish a user's activity, while obfuscating the information about her identity. We retain an identical training and hyperparameter tuning process as explained in Sections 3 and 4, and using the data of 15 users evaluate classifiers' accuracies, with the goal of achieving high activity classification accuracy and low user identity classification.

For the former, we achieve 57% accuracy, which is 2% lower than the performance of the detached activity classifier. More interestingly, however, the encoder is able to properly obfuscate the information related to identity information, as the accuracy of the user identity classifier drops down to 6%, which is on par with the majority classifier, and 73%-points lower than the best-performing classifier constructed without the adversarial obfuscation. Therefore, we conclude that our approach is capable of multidimensional information obfuscation, dependant on a given model.

6 DISCUSSION

6.1 Simultaneous Increase Of User Inference and Activity Obfuscation

One of our initial goals was to limit the loss of user identification accuracy inflicted by the information obfuscation performed by the encoder. To our surprise, the results show that obfuscating the information passed to the user identity classifier *increases* the classification accuracy. A similar observation is noted by the authors of the PAN system [11].

User identification and activity recognition lie in a more orthogonal space than anticipated, meaning that the removal of activity related information boosts performance of the user classifier. As our dataset takes a more traditional approach of hand-crafting the initial feature vector, importance of the encoder is even more prominent, as it does not only filter the data of noise, but also maximises the potential of generated feature space for user classification.

With the learnt insights, we conclude that simultaneously classifying users and providing user privacy are not mutually exclusive traits one should search for a balance between, but rather complementary tools that can frame highly performant authentication systems while retaining user privacy.

6.2 Implications for Real-World Implementations

Modern authentication protocols heavily rely on passwords that are both inconvenient to use and present multiple security risks [15, 20, 25]. One attempt to alleviate these issues is the *OAuth 2.0* protocol that enables delegation of authorisation to third party services, thus lowering the number of different credentials a user has to remember, and consequently increasing security, and reducing the cognitive load incurred during the authentication. We believe that the approach presented in this work has the potential to significantly increase the capabilities of such a protocol.

With the rise of IoT deployments, an increasing number of public and private spaces already provide a plethora of authentication-ready sensor modalities [2], representing solid foundations for the introduction of IoT-supported

authentication. Moreover, with the existing integration of *OAuth 2.0*, the adoption of novel authentication methods requires only minimal effort from application developers – they merely have to issue an API call to an OAuth server. In our case, such a call would replace passwords with behavioural biometrics. Furthermore, training of the complete network presented in this paper can be performed in a privacy retaining way. The utilisation of techniques such as federated learning alleviates the need for the gathered sensor data to leave the local network of IoT devices. Instead, the encoder could be distributed across devices within the household or even corporate network. The data would be transformed by the encoder and then sent to a remote server to train the user and activity classifiers.

6.3 Limitations

To the best of our knowledge, our previous work [9] provide the only publicly available dataset of multi-modal behavioural biometrics collected in an IoT environment that does not focus on mobile devices, but instruments the space instead. Other datasets in literature are created by collecting data from sensors installed on the person itself as in [22]. This dataset, however, is limited with respect to the sensed modalities, number of users and, more importantly, by the number, the granularity, and the quality of the performed activities.

The absolute user identification accuracy numbers we achieve (i.e. 79% accuracy with 15 users), despite representing an almost 20% improvement over the classifiers presented in [9], are still too low for a reliable identification system. As we note, the variation of the (un)familiarity with the prescribed tasks across different runs likely lead to changing behaviour. To avoid this, in future we plan to collect a dataset of users performing multiple runs of more familiar tasks.

On the activity obfuscation aspect, we were successful in confusing the classifier tasked with distinguishing among three rather complex activities. In future we would like to assess whether the high level of obfuscation demonstrated in this paper is still present in case of shorter activities of daily living.

Somehow related is the question of how our solution would perform if the to-be-obfuscated class is heavily dependent on data from a particular sensor. It was previously noted that excluding certain sensing modalities can have a negative effect on user identification in IoT environments [9]. Should our adversarial approach be constructed to obfuscate, for example, the speed at which a person is typing on a shared terminal, we hypothesise that our training algorithm would produce an encoder that would more severely obfuscate data of the accelerometer and gyroscope placed on the keyboard, as those two are crucial in determining a person's typing speed. This, in turn, could have a negative effect on the approach's ability to successfully identify the user. Including a larger number of independent, yet relevant modalities may address this concern.

7 CONCLUSION

In this work we tackled the issue of privacy-preserving behavioural biometric-based user identification in IoT environments. We successfully implemented a three-part privacy-preserving adversarial learning system consisting of the encoder, and user- and activity- (adversarial) classifiers. On a 15-user dataset collected in an IoT environment we were able, to not only match, but improve user classification accuracy by 3% in comparison to a non-adversarial learning classifier, while the accuracy of activity inference dropped by 30% and is on par with the majority classifier. Experiments with varying number of users demonstrate that our approach is likely to scale beyond 15 users available in the given dataset. Finally, by successfully training an inverted classifier that identifies activities, while obfuscating user identities, we demonstrate the flexibility of the proposed approach and its potential applicability to a range of situations where different aspects of user privacy need to be protected.

REFERENCES

[1] Hamed Alqahtani, Manolya Kavakli-Thorne, and Gulshan Kumar. 2019. Applications of Generative Adversarial Networks (GANs): An Updated Review. *Archives of Computational Methods in Engineering* 28, 2 (2019), 525–552. https://doi.org/10.1007/s11831-019-09388-y

[2] Mahmoud Ammar, Giovanni Russello, and Bruno Crispo. 2018. Internet of Things: A survey on the security of IoT frameworks. *Journal of Information Security and Applications* (2018). https://doi.org/10.1016/j.jisa.2017.11.002

[3] Jagmohan Chauhan, Young D. Kwon, Pan Hui, and Cecilia Mascolo. 2020. ContAuth : Continual Learning Framework for Behavioral-based User Authentication. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 4 (2020), 1–19. https://doi.org/10.1145/3432203

[4] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A. Bharath. 2018. Generative Adversarial Networks: An Overview. , 53–65 pages. https://doi.org/10.1109/MSP.2017.2765202

[5] 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* (2013). https://doi.org/10.1109/TIFS.2012.2225048

[6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144. https://doi.org/10.1145/3422622

[7] 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. *2012 IEEE 5th International Conference on Biometrics: Theory, Applications and Systems, BTAS 2012* (2012), 8–15. https://doi.org/10.1109/BTAS.2012.6374552

[8] Atul N. Kataria, Dipak M. Adhyaru, Ankit K. Sharma, and Tanish H. Zaveri. 2013. A survey of automated biometric authentication techniques. *2013 Nirma University International Conference on Engineering, NUiCONE 2013* (2013), 1–6. https://doi.org/10.1109/NUiCONE.2013.6780190

[9] Andraž Krašovec, Daniel Pellarini, Dimitrios Geneiatakis, Gianmarco Baldini, and Veljko Pejović. 2020. Not Quite Yourself Today: Behaviour-Based Continuous Authentication in IoT Environments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 4 (2020).

[10] Fudong Li, Nathan Clarke, Maria Papadaki, and Paul Dowland. 2014. Active authentication for mobile devices utilising behaviour profiling. *International Journal of Information Security* 13, 3 (2014), 229–244. https://doi.org/10.1007/s10207-013-0209-6

[11] Sicong Liu, Junzhao Du, Anshumali Shrivastava, and Lin Zhong. 2019. Privacy adversarial network: Representation learning for mobile data privacy. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 4 (2019). https://doi.org/10.1145/3369816

[12] Ahmed Mahfouz, Tarek M. Mahmoud, and Ahmed Sharaf Eldin. 2017. A survey on behavioral biometric authentication on smartphones. *Journal of Information Security and Applications* 37 (2017), 28–37. https://doi.org/10.1016/j.jisa.2017.10.002

[13] Apostolos Malatras, Dimitris Geneiatakis, and Ioannis Vakalis. 2017. On the efficiency of user identification: a system-based approach. *International Journal of Information Security* 16, 6 (2017), 653–671. https://doi.org/10.1007/s10207-016-0340-2

[14] Philip Marquardt, Arunabh Verma, Henry Carter, and Patrick Traynor. 2011. (sp) iphone: Decoding vibrations from nearby keyboards using mobile phone accelerometers. In *Proceedings of the 18th ACM conference on Computer and communications security*. 551–562.

[15] Robert Morris and Ken Thompson. 1979. Password Security: A Case History. *Commun. ACM* 22, 11 (1979), 594–597. https://doi.org/10.1145/359168.359172

[16] 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* (2016). https://doi.org/10.1109/MSP.2016.2555335

[17] Soumen Roy, Devadatta Sinha, and Utpal Roy. 2017. User authentication: keystroke dynamics with soft biometric features. *Internet of Things (IoT): Technologies, Applications, Challenges and Solutions* (2017), 99.

[18] Zhang Rui and Zheng Yan. 2019. A Survey on Biometric Authentication: Toward Secure and Privacy-Preserving Identification. *IEEE Access* 7 (2019), 5994–6009. https://doi.org/10.1109/ACCESS.2018.2889996

[19] Chao Shen, Shichao Pei, Zhenyu Yang, and Xiaohong Guan. 2015. Input extraction via motion-sensor behavior analysis on smartphones. *Computers and Security* 53 (2015), 143–155. https://doi.org/10.1016/j.cose.2015.06.013

[20] Chao Shen, Tianwen Yu, Haodi Xu, Gengshan Yang, and Xiaohong Guan. 2016. User practice in password security: An empirical study of real-life passwords in the wild. *Computers and Security* 61 (2016), 130–141. https://doi.org/10.1016/j.cose.2016.05.007

[21] Weidong Shi, Jun Yang, Yifei Jiang, Feng Yang, and Yingen Xiong. 2011. SenGuard: Passive user identification on smartphones using multiple sensors. *International Conference on Wireless and Mobile Computing, Networking and Communications* (2011), 141–148. https://doi.org/10.1109/WiMOB.2011.6085412

[22] Zdeňka Sitová, Jaroslav Šeděnka, Qing Yang, Ge Peng, Gang Zhou, Paolo Gasti, and Kiran S Balagani. 2015. HMOG: New behavioral biometric features for continuous authentication of smartphone users. *IEEE Transactions on Information Forensics and Security* 11, 5 (2015), 877–892.

[23] Daniel F. Smith, Arnold Wiliem, and Brian C. Lovell. 2015. Face recognition on consumer devices: Reflections on replay attacks. *IEEE Transactions on Information Forensics and Security* 10, 4 (2015), 736–745. https://doi.org/10.1109/TIFS.2015.2398819

[24] Lina Yao, Quan Z. Sheng, Boualem Benatallah, Schahram Dustdar, Xianzhi Wang, Ali Shemshadi, and Salil S. Kanhere. 2018. WITS: an IoT-endowed computational framework for activity recognition in personalized smart homes. *Computing* 100, 4 (2018), 369–385. https://doi.org/10.1007/s00607-018-0603-z

[25] Verena Zimmermann and Nina Gerber. 2020. The password is dead, long live the password – A laboratory study on user perceptions of authentication schemes. *International Journal of Human Computer Studies* 133, April 2019 (2020), 26–44. https://doi.org/10.1016/j.ijhcs.2019.08.006