
Toward privacy in IoT mobile devices for activity recognition

Théo Jourdan (INSA-Lyon, Inserm, Creatis, CITI)
Antoine Boutet (INSA-Lyon, Inria, CITI)
Carole Frindel (INSA-Lyon, Inserm, Creatis)

Abstract

Recent advances in wireless sensors for personal healthcare allow to recognise human real-time activities with mobile devices. While the analysis of those data-stream can have many benefits from a health point of view, it can also lead to privacy threats by exposing highly sensitive information. In this paper, we propose a privacy-preserving framework for activity recognition. This framework relies on a machine learning technique to efficiently recognise the user activity pattern, useful for personal healthcare monitoring, while limiting the risk of re-identification of users from biometric patterns that characterizes each individual. To achieve that, we rely on a carefully features extraction scheme in both temporal and frequency domain and apply a generalisation-based approach on features leading to re-identify users. We extensively evaluate our framework with a reference dataset: results show an accurate activity recognition (87%) while limiting the re-identification rate (33%). This represents a slightly decrease of utility (9%) against a large privacy improvement (53%) compared to state-of-the-art baselines.

1 Introduction

The emergence of medical Internet of Things (IoT) devices have paved the way for personal healthcare monitoring at home or in hospital environments. These devices record electronic health measurements from a variety of sensors and send these patient data to an application server to be processed and analysed (e.g., to provide services such as motion tracking or vital signs measurement). Due to their nature, collected data from medical IoT devices are highly sensitive. Advances in wireless communication and web technologies facilitate the remote real-time monitoring of such systems [24]. However, the complex workflow of collected medical data multiplies the security and privacy risks all along the life-cycle of the data including the data collection and transmission [23, 3], as well as the processing and the storage [19]. When such medical data can be accessed by an adversary, risks of privacy threats like leakages of sensitive information or user re-identification are very high [14].

In the context of activity recognition through mobile devices, the challenge is to identify data that can preserve the privacy of individuals while still being relevant enough for machine learning tasks [20]. This challenge raises two important questions: 1) Is the collected data protected enough so that no one can misuse it to infer sensitive information or to re-identify the owner? 2) How to assess whether the protected data are still accurate enough for researchers in the health domain? Achieving this balance between data utility and data privacy is an important objective to send secure and reliable data through mobile devices and to strengthen end-user confidence and adoption.

In this paper, we propose a privacy-preserving framework for activity recognition from mobile devices. This framework relies on a machine learning technique to efficiently recognise the user activity pattern, useful for personal healthcare monitoring, while limiting the risk of re-identification of users from biometric patterns that characterizes each individual. To achieve that, firstly we extracted multiple features from raw signal and deeply analysed their impact on both the activity recognition and the

user re-identification. We show that features in temporal domain are useful to discriminate user activity while features in frequency domain lead to discriminate the user identity. Based on this observation, we design a novel privacy-preserving framework. In this framework, data records are processed locally on the user device and only relevant features are extracted. Additionally, features in the frequency domain (i.e., features leading to discriminate users) are normalized. This normalization can be viewed as a generalization-based approach. However compared to other generalization-based approaches based on k -anonymity that are well known to drastically reduce the utility of the protected data [11], our solution keeps a high utility (i.e. activity recognition) while providing a good privacy (i.e. small user re-identification). Once normalized, this information are periodically upload to the application server. Each batch of features is stored independently on the server (i.e., with a different pseudonym) to avoid to link both batches to individuals and batches together. Moreover, to avoid centralizing both the data and the associated identity of their owners on the same node, the mapping between the pseudonyms and the user identities is only retained by the hospital practitioners.

We exhaustively evaluated our framework with a reference dataset. Results show an accurate activity recognition of 87% in average while limiting the user re-identification rate up to 33%. We also compared our solutions against different baselines. Our solution provides a better privacy-utility trade-off with a slightly decrease of utility (9%) against a large increase of privacy (53%).

2 Quantifying activity recognition and user re-identification

We first carried out an extensive evaluation of the capacity to recognise the activity of users and to re-identify them. To do that, we followed a typical methodology for activity recognition from IoT devices (details about our methodology can be found in our companion report [13]). In this work, we used a reference dataset [2]. This dataset is composed of the 3-axial raw data from accelerometer and gyroscope sensors read at a constant frequency of 50 Hz. A group of 30 volunteers were selected to follow a protocol of activities while wearing a smartphone on waist [18]. The experiment was planned in order to contain six basic activities: three static postures (standing, sitting, lying-down) and three ambulation activities (walking, walking-downstairs and walking-upstairs).

Our results (not depicted here for space reason) show that we are able to predict the activity of the user with a very high rate of success. For the **activity recognition** task, our machine learning framework based on Random Forest is able to highly recognise activities with an average accuracy of 0.97 and all the accuracies between 0.94 and 0.99 for the six activities. In addition, we show that without any protection scheme, data from mobile devices act as a personal fingerprint and lead to re-identify users. For the **user re-identification** task, accuracy ranges from 0.82 to 0.96 among the 30 users with an average of 0.90. These results indicate that the data collected from the gesture of users characterizes each individual and can lead to re-identify them with a high success rate. However, the task of re-identification is slightly more difficult than that of recognizing activities with lower accuracy.

Lastly, these experiments were also used to rank features according to their importance. Eight and eleven features were respectively selected for the activity recognition and user re-identification tasks given the correlation and accuracy analysis. Indeed, many features are alike and contain similar information on the original sensor data. Compared to using all 340 features, using only these 19 relevant features lowers only slightly ($< 4\%$) the two classification tasks performance (97% vs 96% for activity classification and 90% vs 86% for user re-identification).

Based on these ranking results, it is interesting to note that the task of activities recognition (i.e., utility) is almost exclusively (9 of the 11 selected features) operated in the time domain whereas the task of user identification (i.e., privacy) is based (5 of the 8 selected features) on features in the frequency domain. These results can be explained by the fact that the activities are mainly distinguished from each other by their level of amplitude in acceleration and gyration and therefore their associated statistics. Conversely, the user identification is more related to the pace or cadence at which this person performs the activity and is strongly related to biomechanics (e.g., age, weight).

3 Privacy-Preserving Activity Recognition Framework

The architecture of our privacy-preserving activity recognition framework is depicted Figure 1. Our privacy-preserving framework involves three premises: the client running on the smartphone of users, the application server storing the features and performing the classification, and the hospital




Figure 1: Architecture of our framework: the user smartphone is leveraged to extract relevant features and only these features are uploaded periodically to the application server.

practitioner monitoring the patient activity. Firstly, IoT devices (e.g. smartwatch) or directly the smartphones perform the data acquisition (❶). In both cases, these raw data are stored locally on the smartphone. The client application then performs the preprocessing, the segmentation and the features extraction following the considered methodology [13]. On the basis of our analysis on the importance of features, this feature extraction only concerns the 19 features identified as important. Moreover, in order to limit the re-identification of users, the client conducts a normalisation scheme which generalises the effect of the different descriptors identified as important for the task of user re-identification. This normalization tries to mitigate their characteristics allowing the re-identification of the user without removing them completely because they also have an impact on the recognition of activities. We distinguished five normalisations (formalised in [13]), each of them referring to the features in the frequency domain identified in the first analysis presented Section 2. Regarding the temporal features, we simply delete them.

As all the aforesaid actions performed on the smartphone only concern the associated user on one batch of data (i.e., one day for instance), the resulting computational cost is cheap. On a commodity computer, these operations applied on all the data of one user spend 2.5 seconds in our experiments. Secondly, the client application associates a random pseudonym to each timestamped batch of features before to periodically upload them to the application server (❷). The client application then sends to the hospital practitioner the list of pseudonyms associated to its identity (❸). When a batch of features is received by the application server, it stores this information in a database (❹). Consequently, each batch in this database does not contain the identity of the user but a random pseudonym. The application server then periodically performs the classification to detect the activity associated to each batch of features. Finally, when the hospital practitioner wants to monitor the activity of a specific users, firstly it retrieves locally all the pseudonyms associated to the specified user and then requests the application server to have the activity history of the specified pseudonyms (❺).

In this work, we assume that the client application and the smartphone on which it is run as well as the server used by the hospital practitioner are trusted. Moreover, we assume that the application server runs on honest but curious public cloud platforms [10]. However, we assume that the adversary is able to collect both part or the entire information stored in the database on the application server and prior knowledge on each user (e.g., from a malicious IoT device) to build a classifier model. This classifier exploits the same methodology than our but with the objective to predict the identity of the user for each batch of data stored in the database. Finally, all communications between the different entities are secured and no information can be inferred from them.

To highlight the benefits of our approach, we compare the performance of our framework with that of two alternatives in term of utility-privacy trade-off. The first alternative follows a perturbation scheme. Similarly to the differentially private approach described in [1] that applies a perturbation scheme in the frequency domain of aggregated time series in the context of location privacy, this alternative (called *perturbation*) adds a Gaussian noise in the signal in frequency domain before the extraction of features. The second alternative is based on the removing of features identified as leading to the user re-identification (called *suppression*).

Figure 2 reports for our solution and the baseline approaches the trade-off between the utility captured by the accuracy to recognise the activity and the privacy captured by the accuracy to re-identify users. For the baseline based on the suppression of features, each point of the curves corresponds to the deletion of a feature (from the 8 selected ones for the re-identification task). For the baseline based on perturbation, in turn, each point refers to the addition of an increasing fixed amount of noise (noise is centered on zero and its standard deviation is, for each point, increased by 2). Finally, in our

framework, each point corresponds to the normalisation of a growing number of features (in order of increasing importance).

Results show that the suppression approach (slope: 0.12) seems the most advantageous in terms of compromise between utility and privacy. However it is very quickly limited by the number of selected features and therefore in privacy and utility metrics; for instance the best obtained performance are respectively 0.66 and 0.93. The perturbation approach (slope: 0.34) is very effective in loss of identification however at the cost of a very important loss of utility too, with for best performance in privacy and utility metrics respectively 0.51 and 0.84. Our approach is between the two (slope: 0.21) and provides the best utility and privacy trade-off (respectively 0.87 and 0.33). Our approach based on normalisation gives a better control on the weight of each feature in the protection, unlike the suppression approach for which limits their impact to consideration or not. Lastly, we also considered an adversary that trains a classifier only with features leading to the re-identification, in this case the accuracy in term of re-identification is less efficient than with our framework (0.17).




Figure 2: Our framework provides a better utility and privacy trade-off than baseline approaches.

4 Related work

With the technological advances of recent years, the medical domain is changing fast raising important security and privacy issues [5, 21]. These concerns emerge at multiple stages in the life-cycle of the data [19, 9, 23, 3]. Although gesture recognition attracts many attention currently [22], to the best of our knowledge, our work is the first one that addresses the protection of data dedicated to activity recognition through wearable devices in the medical domain. The identification of relevant features for both the activity recognition and the user re-identification is also novel. Several well known reported user re-identifications have shown that hiding explicit identity information through pseudonymity is not enough to guarantee the anonymity of users [14]. Indeed, many criteria lead to uniquely identifying users [15, 6, 16, 8]. Following these studies, we also demonstrate in this paper that an user can be easily identified from its gestures collected by sensors. Compared to other approaches that obfuscate independently every record [4], only features leading to the re-identification of users are obfuscated. In addition, although this obfuscation based on a normalization does not provide the same privacy guaranty as other generalization-based approaches ensuring k -anonymity, the utility (i.e., activity recognition) remains high while providing a good privacy (i.e., a small re-identification rate). Lastly, splitting sensitive information (i.e., both the identity of users and their data) on different nodes have already showed its benefits in terms of privacy [12, 17]. In addition, by processing the signals at the edge of the network on the smartphone of users, our framework reduces the operational costs of the application [7] and strengthens the control of users on their data.

5 Conclusion

We present a privacy-preserving IoT framework in the context of activity recognition for healthcare monitoring with wearable devices. Our framework processes the signal and extracts relevant features locally on the user smartphone. In addition, accordingly to the observation that the frequency domain prevails in the user identification task, a normalization is performed on the frequency-based features to obfuscate the re-identification of users. Finally, only a set of features unlinked to the identity of its owner is uploaded to the application server which is then able to recognise the activity of the users with a high accuracy while reducing the risk of user re-identification. An extensive experimental

validation of our framework has been performed on reference data sets yielding good results in terms of privacy-utility trade-off: a high activity recognition with few user re-identification. An interesting research venue for future works will be the formalisation of the privacy guarantee of our solution.

References

- [1] G. Acs and C. Castelluccia. A case study: Privacy preserving release of spatio-temporal density in paris. In *KDD*, pages 1679–1688, 2014.
- [2] D. Anguita, A. Ghio, L. Oneto, X. Parra, and J. L. Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In *ESANN*, 2013.
- [3] D. Aranki and R. Bajcsy. Private disclosure of information in health tele-monitoring. *CoRR*, abs/1504.07313, 2015.
- [4] R. Assam, M. Hassani, and T. Seidl. Differential private trajectory obfuscation. In *MOBIQUITOUS*, pages 139–151, 2013.
- [5] E. Ayday and M. Humbert. Inference attacks against kin genomic privacy. *S&P*, 15(5):29–37, 2017.
- [6] A. Boutet, S. Ben Mokhtar, and V. Primault. Uniqueness Assessment of Human Mobility on Multi-Sensor Datasets. Research report, LIRIS UMR CNRS 5205, Oct. 2016.
- [7] A. Boutet, D. Frey, R. Guerraoui, A.-M. Kermarrec, and R. Patra. Hyrec: Leveraging browsers for scalable recommenders. In *Middleware*, pages 85–96, 2014.
- [8] P. Eckersley. How unique is your web browser? In *PETS’10*, pages 1–18, 2010.
- [9] P. Gard, L. Lalanne, A. Ambourg, D. Rousseau, F. Lesueur, and C. Frindel. A secured smartphone-based architecture for prolonged monitoring of neurological gait. In *HealthyIoT*, pages 3–9, 2018.
- [10] O. Goldreich. Cryptography and cryptographic protocols. *Distrib. Comput.*, 16(2-3):177–199, 2003.
- [11] M. Gramaglia and M. Fiore. Hiding mobile traffic fingerprints with GLOVE. In *CoNEXT*, pages 26:1–26:13, 2015.
- [12] S. Guha, M. Jain, and V. N. Padmanabhan. Koi: A location-privacy platform for smartphone apps. In *NSDI*, pages 183–196.
- [13] T. Jourdan, A. Boutet, and C. Frindel. Toward privacy in IoT mobile devices for activity recognition. In *MobiQuitous 2018*, pages 1–10, 2018.
- [14] L. L. Confidentiality and privacy of electronic medical records. *JAMA*, 285(24):3075–3076, 2001.
- [15] D. Manousakas, C. Mascolo, A. R. Beresford, D. Chan, and N. Sharma. Quantifying privacy loss of human mobility graph topology. *PETS*, 2018(3):5–21, 2018.
- [16] R. Masood, B. Z. H. Zhao, H. J. Asghar, and M. A. Kâafar. Touch and you’re trapp(ck)ed: Quantifying the uniqueness of touch gestures for tracking. *PoPETS*, 2018(2):122–142, 2018.
- [17] A. Petit, T. Cerqueus, S. Ben Mokhtar, L. Brunie, and H. Kosch. PEAS: Private, Efficient and Accurate Web Search. In *TrustCom*, 2015.
- [18] J. L. Reyes-Ortiz. *Smartphone-based human activity recognition*. Springer, 2015.
- [19] M. Rushanan, A. D. Rubin, D. F. Kune, and C. M. Swanson. Sok: Security and privacy in implantable medical devices and body area networks. In *S&P*, pages 524–539, 2014.
- [20] B. Seref and E. Bostanci. Opportunities, threats and future directions in big data for medical wearables. In *BDAW*, pages 15:1–15:5, 2016.

- [21] F. Tramèr, Z. Huang, J.-P. Hubaux, and E. Ayday. Differential privacy with bounded priors: Reconciling utility and privacy in genome-wide association studies. In *CCS*, pages 1286–1297, 2015.
- [22] H. Watanabe, T. Terada, and M. Tsukamoto. Gesture recognition method based on ultrasound propagation in body. In *MOBIQUITOUS*, pages 288–289, 2016.
- [23] D. Wood, N. Apthorpe, and N. Feamster. Cleartext data transmissions in consumer IoT medical devices. In *IoT S&P*, pages 7–12, 2017.
- [24] Z.-K. Zhang, M. C. Y. Cho, C.-W. Wang, C.-W. Hsu, C.-K. Chen, and S. Shieh. Iot security: ongoing challenges and research opportunities. In *SOCA*, pages 230–234, 2014.