Privacy-Preserving Sensor Data Analysis for Edge Computing

Mohammad Malekzadeh
Queen Mary University of London
m.malekzadeh@qmul.ac.uk

Richard G. Clegg
Queen Mary University of London
r.clegg@qmul.ac.uk

Andrea Cavallaro
Queen Mary University of London
a.cavallaro@qmul.ac.uk

Hamed Haddadi
Imperial College London
h.haddadi@imperial.ac.uk

ABSTRACT

To better protect users' privacy while using sensor-equipped devices, we need to move from the current binary setting of granting or not permission to applications for collecting raw sensor data, toward a model that allows each application to get access over a limited range of inferences according to the provided services. In this research, we aim to develop such an integrated sensing, privacy-preserving framework for managing access to sensor time-series data. We have already proposed two novel feature-learning architectures, *Replacement AutoEncoder* (RAE) and *Guardian-Estimator-Neutralizer* (GEN), that locally transform and release personal sensor data in a way that user-defined sensitive information is protected without losing cloud services' utility. The experimental results conducted on activity recognition tasks, using real-world datasets, show that both RAE and GEN can retain the recognition accuracy of state-of-the-art techniques, while concurrently preserving the privacy of sensitive information. In this report, we formulate the underlying utility-privacy-simplicity problem we face and briefly discuss the achieved progress. In ongoing work, we are improving the existing frameworks towards a practical, real-time, and efficient edge computing platform to be deployable in real-world sensing infrastructures.

KEYWORDS

Privacy, Sensor Data, Feature Learning, Time-Series Analysis

1 INTRODUCTION

Recent advances in mobile and ubiquitous computing technologies have accelerated the interest in cloud services. However, omnipresent data gathering and user tracking lead to inherent data security and privacy risks. There is a growing concern about how personal data are used when users grant cloud-based applications direct access to the sensors embedded in smart devices. For instance, time-series data generated by motion sensors, embedded in smartphones, reflect directly users' activities and indirectly their physical and demographic attributes [4]. Although analyzing motion sensor data can enable providing useful applications to the users, such as health and activity monitoring, a lot of sensitive information can be revealed by continuously collecting such high-resolution temporal data.

Difficulties in trusting external data collectors motivate us to investigate for a solution which enables us to transform data before releasing them to a third party. We should move from the current binary setting of granting or not permission to an application, toward a privacy-preserving model that helps users to grant each application permission over a limited range of inferences according to the provided service. To achieve this, we need to set up a transparent data flow to efficiently consider all the existing concerns and assess what action should be taken at each stage (see Figure 1).

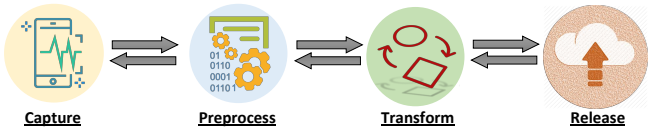

Figure 1: Sensor data flow for privacy-preserving analysis.

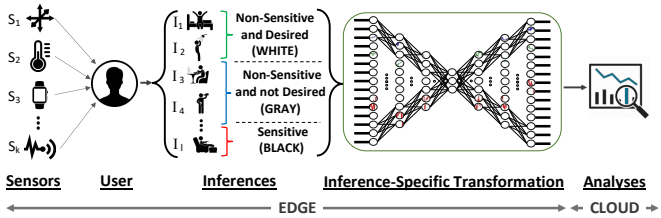

Figure 2: Envisioned Architecture for building a trusted path from Edge to Cloud.

Since after releasing data, it will reside in the cloud side and users no longer have any control over their data, building a good tradeoff between information loss, privacy guarantee, and cost of transformation method is the most challenging task here. The ultimate goal is eliminating the possibility of inferring unwanted sensitive information but preserving that information which users wish to release for getting the desired services from the cloud. In the following, we introduce two data transformation mechanisms based on *deep representation learning* which have been proposed for the users' activity recognition (see Figure 2). Thus, in this report, we first discuss how we can automatically learn discerning features from time-series data to distinguish sensitive and non-sensitive information, then we explain questions and challenges which are still open and need to be addressed in the future work.

2 INFERENCE-SPECIFIC TRANSFORMATION

To make inference on sensor time-series data, which are collected through repeated measurements, there are two possible views: (i) Making *temporal inference*, which means each section of time-series can be assigned to a specific inference that can be *sensitive* to user, *non-sensitive* to user (including inferences that are *desired* for applications; see Figure 3). In this case, users want to gain utility by

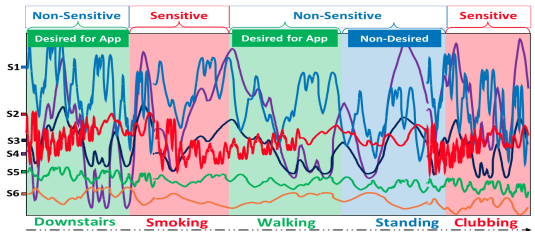

Figure 3: Temporal Inferences: each section of time-series can be assigned to a specific inference that can be *sensitive* to user, *non-sensitive* to user, or *desired* for application.

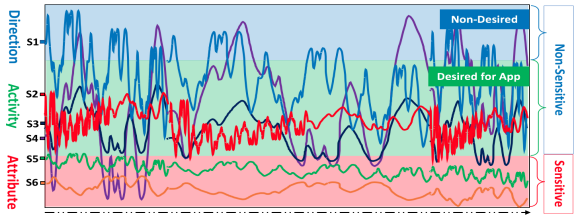

Figure 4: Concurrent Inferences: information available in each section of time-series can be used to concurrently make both *sensitive* and *non-sensitive* inferences.

Activities	Subject	List of Inferences	OF_1	TF_1
0 : null	#1	$I_w = \{1, 2, 3, 4, 9, 10, 11\}$	94.11	90.15
1 : open window		$I_b = \{5, 6, 7, 8\}$	**95.75**	**0.26**
2 : close window		$I_g = \{0\}$	95.04	96.54
4 : water a plant	#2	$I_w = \{2, 3, 5, 6, 7, 9\}$	96.10	92.13
5 : turn book		$I_b = \{4, 8, 10\}$	**96.96**	**0.52**
6 : drink a bottle		$I_g = \{0, 1\}$	95.70	97.56
7 : cut with knife				
8 : chop with knife				
9 : stir in a bowl				
10 : forehand				
11 : backhand				
12 : smash				

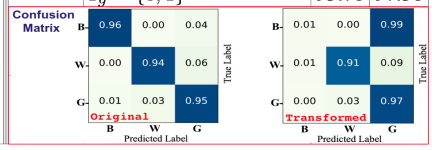

Figure 5: F1-score of original time-series, OF_1, and transformed version, TF_1, for recognizing different types of activities: desired, I_w, sensitive, I_b, and non-sensitive, I_g. Classification confusion matrices shows that after transformation, almost all of the sensitive activities, B, are recognized as non-sensitive one. G, while the recognition of desired activities, W, is kept accurate.

sharing desired sections as well as to protect their privacy by hiding sensitive sections. (ii) Making *concurrent inferences*, which means information available in each section of time-series can be used to make both *sensitive* and *non-sensitive* inferences. Thus, users want to eliminate the possibility of extracting sensitive information, throughout the entire period, while keeping the extraction of non-sensitive information as accurate as possible (see Figure 4).

For this purpose, we have first proposed a replacement approach in order to protect sensitive *temporal inferences* [5]. We introduced *Replacement AutoEncoder* (RAE), a deep feature learning algorithm which learns how to replace discriminative patterns in data, that

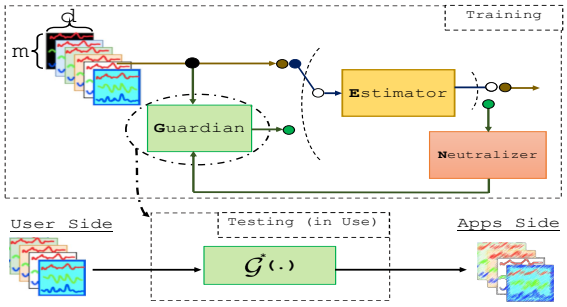

Figure 6: The *Guardian* provides an inference-specific transformation, the *Estimator* guides the Guardian by estimating sensitive and non-sensitive information in the transformed data, and the *Neutralizer* is an optimizer that helps the Guardian converge to a near-optimal transformation function.

Dataset	Inf.	S_d	$\hat{S}_d$
MotionSense	$I_\mathbf{a}$	95.08	93.71
	$I_\mathbf{g}$	95.15	49.32
MobiAct	$I_\mathbf{a}$	94.31	90.46
	$I_\mathbf{g}$	93.74	49.83

Figure 7: Activity recognition, $I_\mathbf{a}$, and gender classification $I_\mathbf{g}$. Accuracy for *original*, S_d, and *transformed*, $\hat{S}_d$, data in percent (%).

correspond to sensitive inferences, with some patterns that have been more observed in non-sensitive inferences, to hide sensitive information. This efficiency is achieved by defining a user-customized objective function for deep autoencoders. Our *replacement* method will not only eliminate the possibility of recognizing sensitive inferences, it also eliminates the possibility of detecting the occurrence of them. That is the main weakness of other approaches such as filtering or randomization. We evaluate the efficacy of the algorithm with activity recognition, using extensive experiments on three benchmark datasets. A sample of results is shown in Figure 5. You can find for more details and results in [5].

Afterward, for protecting *concurrent inferences*, we have designed another learning architecture, called *Guardian-Estimator-Neutralizer* (GEN) [4]. Again, for the specific use-case of activity recognition, we conducted experiments on two other real-world datasets of smartphone's motion sensors (one of them is collected by the authors and is now publicly available[1]). Results indicate the GEN (Figure 6) establishes a good trade-off between application's utility and data subjects' privacy, by maintaining the usefulness of the transformed data for activity recognition (with around an average loss of three percentage points) while almost eliminating the possibility of gender classification (from more than 90% to around 50%, the target random guess). Some results are shown in Figure 7. Read [4] for more explanation and results.

[1]https://github.com/mmalekzadeh/motion-sense

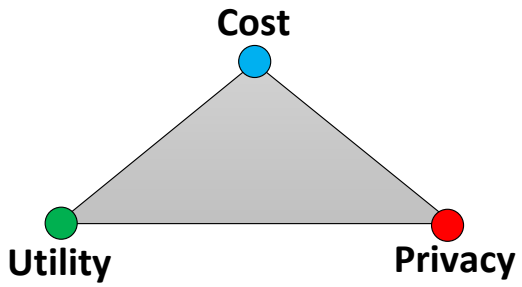

Figure 8: A realistic approach to privacy-preserving sensor data analysis should consider a utility-privacy-cost tradeoff.

3 OPEN QUESTIONS

Protecting privacy in sensor data analysis is a very challenging task, especially when data are gathered from different sources. A principal task is defining proper measures that can truly assess the privacy-utility-cost tradeoff (see Figure 8). Most of the well-known measures focus on protecting the identity of a participant without considering the information content of the corresponding data. An acceptable solution for sensor data needs to accurately measure the amount of original information that is still contained in the transformed data as well as the cost and complexity we incur for achieving this.

Local Differential Privacy [1] offers a strong privacy guarantee for access to individual data controlled by a parameter ϵ, called the privacy loss. Although, the privacy guarantee provided by locally differentially private mechanisms can be very strong for a single round of sensor data release, but it degrades rapidly when data is released regularly. It can be mitigated by adding more noise at each round, but it will also exterminate the utility of the data in the long-term data release. Thus, we need to look for a proper measure of privacy guarantee in future work.

Finally, we should provide a *privacy-by-design* solution that preserves the privacy of sensitive information without the need to trust third parties and simultaneously allow data subject to benefit from cloud-based services for their desired applications. Figure 9 shows the high-level architecture of the envisioned Mediator; *a privacy-preserving platform operates between cloud-based services and data subject.* The Mediator is supposed to be a trusted application which will operate at the Edge side and under user's control. Generally, the main goal of the Mediator is to prevent the reconstruction of original time-series from transformed ones, yet allow to accurately estimate some statistics despite the transformation. The computation cost of curating and mediating time-series data should be considered in proposing any method for the Mediator.

4 MOTIONSENSE DATASET

Unlike other areas of research, it is pretty hard to find a dataset of sensory data which includes information about both activities and personalities of the data subject. For this reason, we have been collecting a dataset called MotionSense [2], as part of databox project [2]. This dataset includes time-series of accelerometer acceleration and

[2]https://github.com/mmalekzadeh/motion-sense

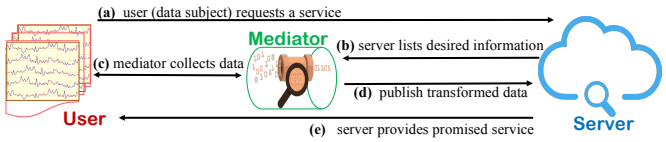

Figure 9: High-Level Architecture: a privacy-preserving mediator between cloud-based servers and data subject. Sensitive sections of personal time-series are transformed with the Mediator, based on the relevant application, and then shared with the third party's server.

gravity) , gyroscope, and attitude (pitch, roll, yaw) data collected by an iPhone 6s kept in the participant's front pocket using SensingKit [3]. There are 24 participants in different range of age, weight, height and gender who performed 6 type of activities in 15 different trials: downstairs, upstairs, walking, jogging, siting, and standing. All of the participant followed the same scenario in a same environment and condition. With this dataset, we aim to look for personal attributes fingerprints in time-series of sensor data. We discuss that there are some attribute-specific patterns in this data which can be used to infer about the data subjects personality as well as their activities.

5 CONCLUSIONS AND FUTURE DIRECTIONS

In this research, we focus on the users' privacy and consider time-series generated by sensors embedded into mobile and wearable devices. We assume these time-series may contain some patterns and features which can reveal a lot of sensitive information about data subject's personal life. We believe that the approach of learning privacy-related features from time-series data will enable us to develop efficient methods for data transformation and help us to enrich existing IoT platform with a robust privacy-preserving time-series data analytics component. Thus, in the future, we will focus on developing the mediator framework by letting users dynamically define their personal privacy policies and inferences or resources the mediator allows applications to access. We will try to better understand the long-term dependencies and discriminative features of time-series, as well as utility-privacy-cost tradeoffs of running proposed methods on edge devices.

REFERENCES

[1] Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. *Journal of Privacy and Confidentiality 7*, 3 (2017), 17–51.

[2] Haddadi, H., Howard, H., Chaudhry, A., Crowcroft, J., Madhavapeddy, A., McAuley, D., and Mortier, R. Personal data: thinking inside the box. In *Proceedings of The Fifth Decennial Aarhus Conference on Critical Alternatives* (2015), Aarhus University Press, pp. 29–32.

[3] Katevas, K., Haddadi, H., and Tokarchuk, L. Sensingkit: Evaluating the sensor power consumption in ios devices. In *12th International Conference on Intelligent Environments (IE'16)* (2016).

[4] Malekzadeh, M., Clegg, R. G., Cavallaro, A., and Haddadi, H. Protecting sensory data against sensitive inferences. *In W-P2DS'18: 1st Workshop on Privacy by Design in Distributed Systems , April 23-26, 2018, Porto, Portugal. ACM, New York, NY, USA, 6 pages. https: //doi.org/10.1145/3195258.3195260.*

[5] Malekzadeh, M., Clegg, R. G., and Haddadi, H. Replacement autoencoder: A privacy-preserving algorithm for sensory data analysis. *The 3rd ACM/IEEE International Conference on Internet-of-Things Design and Implementation, Orlando, Florida, U.S.A.* (17-20 April, 2018).