



DYSAN: Dynamically sanitizing motion sensor data against sensitive inferences through adversarial networks

Antoine Boutet, Carole Frindel, Sébastien Gambs, Théo Jourdan, Rosin
Claude Ngueveu

► To cite this version:

Antoine Boutet, Carole Frindel, Sébastien Gambs, Théo Jourdan, Rosin Claude Ngueveu. DYSAN: Dynamically sanitizing motion sensor data against sensitive inferences through adversarial networks. ACM ASIACCS 2021 - 16th ACM ASIA Conference on Computer and Communications Security, Jun 2021, Hong Kong (Virtuel), China. 10.1145/3433210.3453095 . hal-02512640v3

HAL Id: hal-02512640

<https://inria.hal.science/hal-02512640v3>

Submitted on 24 Jan 2022

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.

DySAN: Dynamically Sanitizing Motion Sensor Data Against Sensitive Inferences Through Adversarial Networks

Antoine Boutet
antoine.boutet@insa-lyon.fr
Univ Lyon, INSA Lyon, Inria, CITI
Villeurbanne, France

Carole Frindel
carole.frindel@insa-lyon.fr
Univ Lyon, INSA Lyon, CNRS, Inserm,
CREATIS UMR 5220, U1206
Villeurbanne, France

Sébastien Gambs
gambs.sebastien@uqam.ca
Université du Québec à Montréal
Montréal, Canada

Théo Jourdan
theo.jourdan@insa-lyon.fr
Univ Lyon, INSA Lyon, CNRS, Inserm,
CREATIS UMR 5220, U1206
Villeurbanne, France

Rosin Claude Ngueveu
ngueveu.rosin-
claudio@courrier.uqam.ca
Université du Québec à Montréal
Montréal, Canada

ABSTRACT

With the widespread development of the quantified-self movement, an increasing number of users rely on mobile applications to monitor their physical activity through their smartphones. However, granting applications a direct access to sensor data exposes users to privacy risks. In particular, motion sensor data are usually transmitted to analytics applications hosted in the cloud, which leverages on machine learning models to provide feedback on their activity status to users. In this setting, nothing prevents the service provider to infer private and sensitive information about a user such as health or demographic attributes. To address this issue, we propose DySAN, a privacy-preserving framework to sanitize motion sensor data against unwanted sensitive inferences (*i.e.*, improving privacy) while limiting the loss of accuracy on the physical activity monitoring (*i.e.*, maintaining data utility). Our approach is inspired from the framework of Generative Adversarial Networks to sanitize the sensor data for the purpose of ensuring a good trade-off between utility and privacy. More precisely, by learning in a competitive manner several networks, DySAN is able to build models that sanitize motion data against inferences on a specified sensitive attribute (*e.g.*, gender) while maintaining an accurate activity recognition. DySAN builds various sanitizing models, characterized by different sets of hyperparameters in the global loss function, to propose a transfer learning scheme over time by dynamically selecting the model which provides the best utility and privacy trade-off according to the incoming data. Experiments conducted on real datasets demonstrate that DySAN can drastically limit the gender inference up to 41% (from 98% with raw data to 57% with sanitized data) while only reducing the accuracy of activity recognition by 3% (from 95% with raw data to 92% with sanitized data).

KEYWORDS

Privacy, Activity Detection, GAN, Utility-Privacy Trade-off

ACM Reference Format:

Antoine Boutet, Carole Frindel, Sébastien Gambs, Théo Jourdan, and Rosin Claude Ngueveu. 2022. DySAN: Dynamically Sanitizing Motion Sensor Data Against Sensitive Inferences Through Adversarial Networks. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

1 INTRODUCTION

The integration of motion sensors in smartphones and wearables has been accompanied by the growth of quantified-self movement [36]. For instance nowadays, users increasingly exploit these devices to monitor their physical activities. Usually, the motion sensor data is not analyzed directly on the device but are rather transmitted to analytics applications hosted in the cloud. These analytics applications leverage machine learning models to compute statistical indicators related to the status of users that are sent back to them. While these analyses can bring many benefits from the health perspective [11, 24, 26], they can also lead to privacy breaches by exposing personal information regarding the individual concerned. Indeed, a large range of inferences can be done from motion sensor data including sensitive ones such as demographic and health-related attributes [12, 15, 16].

Consider for instance the scenario in which Alice, a woman, uses a fitness application on her smartphone to monitor her physical activity. The application performs the activity recognition as well as the activity monitoring in the cloud. However even if the service provider declares it explicitly in the terms of conditions of the service, Alice has no formal guarantees that her data will not be processed to infer other information about her (*e.g.*, for targeting or marketing purposes). Another possible scenario is related to the new trend of insurance companies that propose discount to clients if they accept to use a connected device to follow their daily activity [33]. These data can be used to provide a personalized coaching for better health management but also for early detection of a pathology, which can negatively impact the insurance cost or lead to other types of discrimination. To address the issues raised by these two scenarios, in this work we propose a solution sanitizing the motion sensor data in such a way that it hides sensitive attributes while preserving the activity information contained in the data.

To achieve this objective, we design DySAN, inspired from the framework of Generative Adversarial Networks (GANs) [23] to sanitize the sensor data. More precisely, by learning in a competitive manner several networks, DySAN is able to build models for sanitizing motion data to prevent inferences on a specified sensitive attribute while maintaining a high level of activity recognition. One of the objective of DySAN is also to limit the distortion between the

raw and sanitized data, thus also maintaining a high level of utility with respect to other analysis tasks related to activity monitoring (e.g., steps counting). DySAN has thus to learn sanitizing models in order to find the best trade-off to cope with these conflicting optimizing goals.

Furthermore, our approach aims at addressing the heterogeneous aspect of data collected by motion sensors. Indeed, these sensor data are user dependent and inherently reflect the way each user moves, to the characteristics of the sensors used for data collection and to the evolution of activity during the day. Thus, one unique sanitizing model cannot cope with the heterogeneity of data and provide the best utility/privacy trade-off for all users over time. To solve this issue, DySAN builds a set of diverse sanitizing models by exploring different combination of hyperparameters leading to different balance in privacy protection with respect to sensitive inference, utility preservation in terms of the loss induced for activity recognition, and the data distortion. By doing so, DySAN is able to assess the trained sanitized models and to dynamically select the model providing the best trade-off over time according to the incoming sensor data.

The evaluation of DySAN on real datasets, in which the *gender* is considered as the sensitive information to hide due to the possible risk of discrimination, demonstrates that DySAN can drastically limit the gender inference up to 41% while only inducing a drop of 3% on the accuracy of activity recognition. In addition to preserve activity recognition, DySAN, by limiting data distortion, also preserves the sensor data utility for other analytical tasks such as estimating the number of steps. Moreover, we show that the dynamic model selection of DySAN successfully provides an adaptation of the sanitization according to the incoming user data. This dynamic model selection is especially useful to generalize the sanitization capacity learnt from the dataset used to build the sanitizer models to another dataset with new users with potentially different behaviours. Our dynamic sanitizing method overcomes several shortcomings of the state-of-the-art approaches, namely the use of the same sanitizing model for all users over time, which may lead to a poor privacy-utility trade-off for atypical users. Lastly, we evaluate the cost of operating DySAN on a smartphone and show that the introduced overhead is compatible with real-time processing and that the energy consumption remains reasonable. Our implementation of DySAN as well as the datasets used to assess its performances are publicly available ¹.

The outline of the paper is as follows. First, the problem definition and the system model considered are described in Section 2. Afterwards, DySAN is presented in Section 3 before reviewing the experimental setting as well as the results obtained, respectively in Section 4 and Section 5. Finally, the related work is discussed in Section 6 before concluding in Section 7.

2 PROBLEM DEFINITION AND SYSTEM MODEL

2.1 Overview and system model

We consider a mobile application installed on the user’s smartphone aiming to monitor its physical activity. The smartphone of the user

is assumed to be under the control of the user and thus trusted while the service provider responsible for activity monitoring is not. In the “classical” (i.e., non-private) version of activity monitoring, the data is acquired by sensors (e.g., accelerometer, gyroscope), retrieved by a dedicated mobile application and then send to a server hosted in the cloud (i.e., no classification is made locally). This server leverages machine learning models to identify the activity of the user or to estimate other physical activity features (e.g., number of steps). The server is considered to follow the honest-but-curious adversary model in the sense that it may also try to infer additional sensitive information from the sensor data.

For the rest of the paper, we consider the *gender* as being the sensitive attribute to protect as it could lead to risk of discrimination. In addition, the gender could be inferred from the list of performed activities and their associated frequencies in case of unbalanced data distribution between men and women (which is not the case in the datasets considered in this paper). However, our approach is much more generic and could be applied to protect other sensitive attributes (e.g., handicap or race).

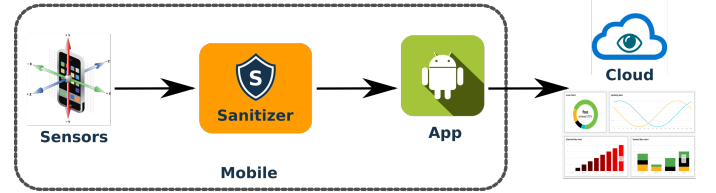


Figure 1: DySAN locally sanitizes the motion sensor data on the smartphone to prevent the cloud-based service from inferring an unwanted sensitive attribute while allowing this service to detect the activity performed by the users as well as compute statistics related to their physical activity.

An overview of DySAN is shown in Figure 1. To avoid an unwanted exploitation of the motion sensor data, these data are sanitized by DySAN before being transmitted to the mobile application. This sanitizing process removes the correlations with the sensitive attribute in the sensor data while preserving the information necessary to detect the activity performed by a user. In addition, DySAN also aims at limiting the distortion between the raw and sanitized data to preserve the utility for other analytical tasks. Finally, the resulting sanitized data are sent to an analytics application hosted in the cloud, exploiting machine learning models to classify the users activity and compute statistics related to their physical activity.

Ideally to limit the privacy risks, the sanitizing step should be done as early as possible to prevent other applications to have access to the raw sensor information. For instance, we consider that DySAN could be deployed in the trusted environment of the smartphone to prevent the mobile application to have a direct access to the sensor data but only from the output of DySAN (thus ensuring that the mobile application uses only sanitized data). Afterwards, the mobile application sends the sanitized data to a server hosted in the cloud similarly to the non-private scenario.

2.2 Problem statement

More formally, we consider raw motion sensor data (denoted by A) captured through accelerometer and gyroscope that sample 3-axial

¹DySAN: <https://github.com/DynamicSanitizer/DySAN>

signals with a frequency of 50 Hz. To enable activity recognition over time, the raw sensor data are split in sliding windows, in which each sliding window is considered to be a sample of a single activity (*i.e.*, by assumption the user cannot perform two different activities during a single sliding window). The choice of the window size is not trivial, especially for an activity recognition task and has to be well calibrated. Indeed, a small window size could split an activity signal while large window size could contain multiple activity signals. Knowing that in average the walking pace is not less than 1.5 steps per second [4], the window length T is chosen to be 2.5 seconds with an overlap of 50 % to match that of a walking cycle of two steps.

We assume a population of N users contained in a dataset X storing all user's data. This dataset includes the raw sensor data as well as the label associated with the activity performed by the user (denoted by a multi-valued attribute Y), the binary sensitive attribute (denoted by S) and a timestamp. Thus, the dataset $X = \{A, Y, S\}$ in which $A = (A_1, \dots, A_T)$.

The objective of DySAN is to protect the user motion sensor data against sensitive attribute inferences while maintaining data utility. More formally, we aim at learning a set of sanitizers $San_{\alpha, \lambda, \beta}$ for various values of the hyperparameters α , λ and β . Each sanitizer will transform the original data X into $\bar{X} = San_{\alpha, \lambda, \beta}(X) = \{\bar{A}, Y, S\}$; $\bar{A} = (\bar{A}_1, \dots, \bar{A}_T)$. This set of sanitizers is learned so that it is difficult to build a discriminator D_{isc} trained to predict S from the sanitized data and activities $\{\bar{A}Y\}$ while an activity predictor P_{red} trained on the same sensor data (\bar{A}) is able to maintain an accuracy close to the original one (more details in the following section). To preserve further the utility of \bar{X} , the sanitizing process is constraint to minimize the distortion between the original and sanitized data.

Furthermore, DySAN aims to dynamically adapt over time the hyperparameters of the model according to the incoming data of each user. Indeed, while a particular model could provide the best utility/privacy trade-off on average for all users with respect to the training dataset, the model leading to the best trade-off can change when testing on new user (*e.g.*, when the new user data does not fit the data distribution of the training dataset). More formally, the objective is to find for each window of data the sanitizer $\widehat{San}_{\alpha, \lambda, \beta}$ providing the best utility/privacy trade-off for the current incoming data. This trade-off is defined by a metric combining the accuracy of the activity recognition and the inference of the sensitive attribute.

3 DYSAN: DYNAMIC SANITIZER

Before exploiting DySAN, multiple sanitizers corresponding to various utility and privacy trade-offs are built during the offline training phase. These models are then deployed on the smartphone. During the online phase, DySAN dynamically selects on the smartphone the best sanitizer for the incoming sensor data. Both the training and the online phases are summarized in Figure 2 and explained in the following subsections.

3.1 Building multiple sanitizers

The offline training phase is performed only once and aims at learning multiple sanitizers. This training is performed with a reference dataset used in activity recognition, the MotionSense dataset that we describe in Section 4.1. As shown in Figure 2, DySAN is composed

of multiple building blocks that we detail hereafter: 1) a sanitizer, 2) a discriminator, 3) a predictor, 4) a distortion measurement and 5) a multi-objective loss function.

- **Discriminator:** The discriminator D_{isc} guides the sanitizer through the process of removing information related to the sensitive attribute $S \in \{0, 1\}$. In practice, we use a Convolutional Neural Network (CNN), which is well-suited to capture time-invariant features in time series [13]. The architecture of this CNN is presented in Appendix A.1. The training of the discriminator is based on a loss function measuring the Balanced Error Rate BER [10] between the output of the discriminator and the ground truth sensitive attribute, which is defined as:

$$BER(D_{isc}(\bar{A}, Y), s) = \frac{1}{2} \left(\sum_{s=0}^1 P(Disc(\bar{A}, Y) \neq s | S = s) \right). \quad (1)$$

The value of BER ranges between 0 and 0.5, in which a value close to 0 corresponds to a perfect accuracy for the prediction of the sensitive attribute while 0.5 means the discriminator is unable to retrieve any information about the sensitive attribute from the sanitized data. Hereafter, we will refer to this loss by $Loss_{Sensitive}$.

- **Predictor:** The predictor P_{red} aims at helping the sanitizer in preserving as much information as possible with respect to the activity recognition task. We also use a CNN for the predictor that has been optimized for predicting the user activity from the sanitized data. The architecture of this CNN is presented in Appendix A.2. Thus, the predictor is trained to maximize the accuracy in inferring activities from the output of the sanitizer. We also use the balanced error rate as the loss function that should minimize the error between the output of the predictor and the ground truth of the activity: $BER(P_{red}(\bar{A}), y)$. For the rest of the paper, we will refer to the predictor loss as $Loss_{Activities}$.
- **Distortion measurement:** The last constraint on the sanitizer is the minimization of data distortion between the raw and sanitized data. Specifically, this distortion should be limited to keep as much information as possible in the sensor data for subsequent analytical tasks. The data distortion is measured through the L_1 loss function denoted $l1$, applied independently on each attribute. For two vectors A_i and \bar{A}_i , corresponding respectively to the raw and sanitized sensor data, the loss function is defined as follows:

$$l1(A_i, \bar{A}_i) = \frac{1}{N_A} \sum_{j=1}^{N_A} |a_{ij} - \bar{a}_{ij}|, \quad (2)$$

in which N_A is the number of possible values for a particular attribute (*e.g.*, the number of axes of the accelerometer or the gyroscope), $a_{ij} \in A_i$ and i denotes a single observation in the window of length T .

- **Sanitizer:** The sanitizer San modifies the raw data taken as input to remove information correlated with the sensitive attribute while maintaining useful information for activity detection. Since the raw and sanitized data belong to the same space, we have implemented the sanitizer as an auto-encoder. In a nutshell, an auto-encoder is a neural network performing a dimension reduction of the signal to compress

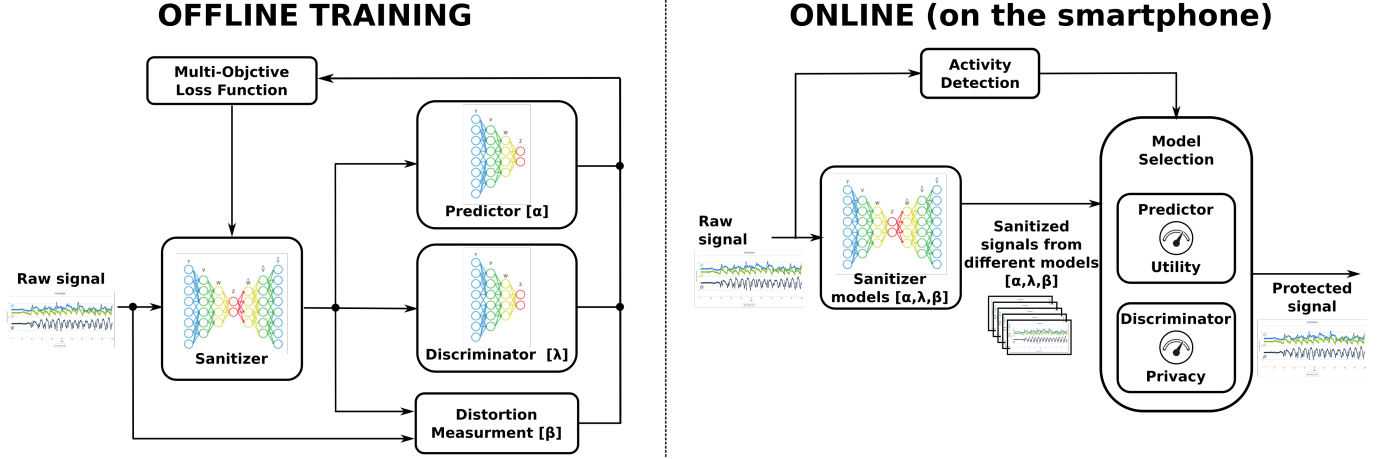


Figure 2: DySAN is composed of two phases: an offline training phase (left) and an online phase (right). The training phase is performed only once and aims to build different sanitizer models that are distinguished by their hyperparameters. Once these sanitizer models deployed on the smartphone, the online phase aims to dynamically choose among these models the most adapted one for each batch of incoming data.

information before trying to reconstruct the input. The sanitizer takes into account the feedback of the discriminator, predictor and distortion measurement to output the sanitized version of the input raw data. More precisely, these different feedbacks are integrated into a multi-objective loss function that should be minimized. The architecture of the auto-encoder is given in Appendix A.3.

- **Multi-objective loss function** The multi-objective loss function J^{San} drives the transformation performed by the auto-encoder to generate the sanitized data \tilde{X} . This loss function takes into account three components, the capacity to detect the activity of the user (*i.e.*, the output of the predictor), the capacity to detect the sensitive attribute (*i.e.*, the output of the discriminator), and the level of distortion introduced in the sanitized data compared to the original one. More formally, the multi-objective is defined as follows:

$$J^{San}(X, San, Disc, Pred) = \{\alpha * d_s(S, Disc(San(X))), \lambda * d_p(Y, Pred(San(X))), \beta * d_r(X, San(X))\}, \quad (3)$$

in which $d_s(x) = \frac{1}{2} - Loss_{Sensitive}$, $d_p = Loss_{Activities}$ and $d_r = \{I1(a_{:,j}, \tilde{a}_{:,j}), \dots\}$ with $a_{:,j}$ representing a dimension of all timesteps of a single sliding window. The term $\frac{1}{2}$ in $d_s(x)$ comes from the objective of maximizing the error of the discriminator, since the sanitizer aims at modifying the data so that the discriminator is no more able to infer sensitive information.

A gradient descent is applied on J^{San} to minimize the global loss function following a similar approach as in [2]. Note that each loss term is weighted with a hyperparameter. More precisely, d_s , d_p and d_r are weighted respectively with α , λ and β . The parameter α represents the relative importance given to the privacy while λ controls the utility (*i.e.*, the quality of activity detection). As we impose the constraint

that $\alpha + \lambda + \beta = 1$, we only adjust α and λ hyperparameters, leaving $\beta = 1 - (\alpha + \lambda)$.

3.2 Training Phase

During the training phase, we build a sanitizer for each set of possible values for the hyperparameters α and λ to explore the domain of the multi-objective loss function. This exploration will allow DySAN to select the best model for each user during the online phase. The training procedure is summarized in Algorithm 1.

In order to optimize the utility and privacy trade-off for a specific set of α and λ (line 1, Algorithm 1), the three neural networks are trained in an adversarial manner. This adversarial training can be seen as a game between the sanitizer on one side and the predictor and the discriminator on the other side. These neural networks compete against each other with opposing objectives until an equilibrium is reached. More precisely, the sanitizer is trained to fool the discriminator and maintained a high activity detection quantified with the predictor while limiting the data distortion. We follow the standard training procedure of GANs consisting in alternating in an iterative manner (at each batch of data) the training of each model with their respective loss function until convergence or until a maximum number of epoch (*i.e.*, we do not consider Competitive Gradient Descent [31]).

Specifically, after initialization (lines 1 – 8) the training of the sanitizer starts with J^{San} while the discriminator and the predictor are frozen (lines 11 – 12). Once the training of the sanitizer has converged, the predictor and the discriminator are trained independently with their respective loss function while the sanitizer is frozen (lines 13 – 20). These two neural networks are trained until convergence (*i.e.*, until the loss no longer decreased) or if a maximum number of iterations, respectively K_{pred} and K_{disc} , is reached. This two-steps process is performed iteratively until an equilibrium is reached.

Algorithm 1 DySAN training algorithm

```
1: Input:  $X, \lambda, \alpha, \max\_epoch, \text{batch\_size}, K_{pred}, K_{disc}$ .
2: Outputs:  $S_{an}, Disc, P_{red}$ .
3: train(M, **trParams): Train the model M using trParams.
4: freeze(M): Freeze the model M parameters and avoid modifications.
5: {Initialisation}
6:  $S_{an}, Disc, P_{red}, X_d = \text{shuffle}(X), X_p = \text{shuffle}(X)$ 
7:  $\text{Iterations} = \frac{|D|}{\text{batch\_size}}$ 
8: {Training Procedure}
9: for  $e = 1$  to  $\max\_epoch$  do
10:   for  $i = 1$  to  $\text{Iterations}$  do
11:     Sample batch  $B$  of size  $\text{batch\_size}$  from  $X$ 
12:      $\text{train}(S_{an}, B, J^{San}, \alpha, \lambda, \text{freeze}(P_{red}), \text{freeze}(Disc))$ 
13:     for  $k = 1$  to  $K_{pred}$  do
14:       Sample batch  $B$  of size  $\text{batch\_size}$  from  $X_p$ 
15:        $\text{train}(P_{red}, B, \text{Loss}_{Activities}, \text{freeze}(S_{an}))$ 
16:     end for
17:     for  $k = 1$  to  $K_{disc}$  do
18:       Sample batch  $B$  of size  $\text{batch\_size}$  from  $X_d$ 
19:        $\text{train}(Disc, B, \text{Loss}_{Sensitive}, \text{freeze}(S_{an}))$ 
20:     end for
21:   end for
22: end for
```

3.3 Online Phase

Once deployed on the smartphone, DySAN is composed of four components as depicted in Figure 2: the sanitizer, the discriminator, the predictor and an activity detection component. Specifically, DySAN knows all the sanitizer, predictor and discriminator models built during the training phase. This set of models correspond to the different possible utility and privacy trade-offs (*i.e.*, set of values explored for the α and λ hyperparameters). The selection of the model is performed by maximizing $S(P, U) = xU + yP$, in which x and y are positive weight coefficients with $x+y = 1$, U the evaluation of the activity done by the predictor, and P the accuracy in terms of privacy as $P = 1 - |0.5 - p|$, in which P is the evaluation of the gender done by the discriminator. Consequently, P is higher when the evaluation of the gender accuracy corresponds to a random guess (*i.e.*, an accuracy of 0.5). According to the expected utility and privacy trade-off, the coefficients x and y can be tuned (the impact of these parameters is assessed in Appendix D).

To find the best sanitizer over time (according to coefficients x and y), DySAN evaluates the utility and the privacy of all models to select the best one. This evaluation requires to know the actual activity performed by the user and the sensitive attribute. While the sensitive attribute can be given by the user, the motion sensor data are not labeled with the activities as it is rather the objective of the activity recognition task to perform this inference.

We use the activity detection component (see Figure 2) to annotate some motion sensor data with their activities on the smartphone. More precisely, we ask the user to follow a specific calibration process at the installation of DySAN. During this process, the user is asked to perform a series of different activities for short

periods to learn a specific supervised classifier to detect his activities. As the quantity of data available to train this classifier is limited, we rely on the use of random forests that are adapted to this context [14]. This random forest (RF) classifier is then used to label the raw data in order to evaluate the utility of all sanitizers. This evaluation is performed on a regular basis (*e.g.*, each period of p windows) and we compute the average accuracy over this basis. By following this process, DySAN is able to identify over time the sanitizer providing the best utility and privacy trade-off defined as a measure combining the accuracy of the activity recognition and the inference of the sensitive attribute.

4 EXPERIMENTAL SETTING

4.1 Datasets

We used two real datasets, which are both publicly available and heavily used in the literature: MotionSense and MobiAct. These datasets contains motion sensor data of subjects doing cyclo- stationary activities (*i.e.*, based on step pattern).

- **MotionSense** [29] contains data captured from an accelerometer (*i.e.*, acceleration and gravity) and gyroscope at a constant frequency of 50Hz collected with an iPhone 6s kept in the front pocket. Overall, a total of 24 participants have performed six activities (*i.e.*, going downstairs, going upstairs, walking, jogging, sitting and standing) during 15 trials in the same environment and conditions.
- **MobiAct** [35] records the data from 58 subjects during more than 2500 trials, all captured with a smartphone in a pocket. This dataset includes signals recorded from the accelerometer and gyroscope of a Samsung Galaxy S3 smartphone with subjects performing nine different types of activities of daily living. For our experiments, we only used the trials corresponding to the same activities as MotionSense.

Both datasets are balanced and contains an equal number of males and females. However, the walking activity is more represented compared to others (Section 5.3). Each activity is performed equally by all subjects in both datasets, making any correlation between the gender and the activity impossible. The datasets are split between training and testing, with 2/3 of trials used for training and validation and 1/3 for testing. These two datasets share similar characteristics, which allows to test the transferability of the models from one dataset to the other. More precisely, the models learned on one dataset can be used to sanitize data from the other dataset. This evaluation corresponds to a more realistic use case and to the best of our knowledge was never considered in previous work related to the sanitization of sensor data.

4.2 Baselines

To assess the performance of DySAN, we considered a set of baselines that we detail hereafter. One of these baselines is based on a random forest classifier [14] while the others are based on GANs [18, 19, 28]. Regarding GAN approaches, authors use an architecture of neural networks slightly different from ours. To provide a fair comparison, we propose to implement their functionalities in our architecture (number of layers, type of CNN, ...). This methodology allows to assess the main characteristics adopted in the baselines

without depending on their choice of architecture that can also have an impact on performance.

- **ORF:** To limit the exposure of the data, in [14] the raw data is preprocessed on the user’s smartphone and only relevant features are transmitted to the application hosted in the cloud. The relevant features are first identified according to the target application (e.g., activity recognition) and selected either in the temporal or the frequency domain. Originally proposed to avoid users re-identification, we adapt this approach to prevent the inference of the sensitive attribute, namely gender. More specifically, we first detect the features that are the most correlated with the gender before normalizing the features in the frequency domain and removing the features in the temporal domain that are not used for the activity classification.
- **GEN:** Similarly to DySAN, GEN (Guardian Estimator Neutralizer) [18] also relies on an adversarial approach to optimize the utility and privacy trade-off. However, this solution does not follow the standard iterative training procedure of GANs as described in Section 3.2. More precisely, the first network, a classifier, is learned once on the raw data to identify both sensitive (e.g., the gender) and non-sensitive information (e.g., the activity). Then the second network, an auto-encoder, is also trained only once through a loss function that does not take into account the data distortion. Finally, the model used in the online phase is the same for all users and corresponds to the best set of hyperparameters identified during the training phase. While this solution relies on a neural network architecture slightly different from ours, we implement GEN by using our architecture. However, to evaluate the performance of GEN in a context of transfer learning, we also use their original neural networks (learned on MotionSense²) to assess its performance on MobiAct.
- **Olympus:** This approach [28] is similar to GEN with the exception that two different neural networks are used to learn the sensitive attributes and to learn non sensitive information. In addition, these classifiers are trained using sanitized data by following an iterative process similar to DySAN described in Section 3.2. However, the loss function does not account for data distortion and the model deployed is the same for all users (i.e., only one sanitizing model trained and used in the online phase). While this approach is used for a different objective (i.e., to avoid users re-identification), we adapt it for activity recognition by using our architecture.
- **MSDA:** This solution [19] can be viewed as an evolution of Olympus in which the loss function driving the training of the auto-encoder accounts for data distortion. However, the model used in the online phase is still the same for all users. While this approach was originally developed with a different purpose in mind (i.e., to avoid re-identification), we adapt this solution for activity recognition by using our architecture. This baseline is the closest to DySAN but without the training of multiple sets of hyperparameters and the dynamic sanitizing model selection in the online phase.

As described in Section 3.3, in the online phase DySAN selects the sanitizing model which provides the best utility and privacy trade-off controlled by parameters x and y . The reported results in the evaluation correspond to an unbalanced trade-off to improve the privacy (i.e., $x = 0.1$ and $y = 0.9$). The impact of different parameters is assessed in Appendix D.

4.3 Evaluation metrics

We evaluated DySAN along both utility and privacy metrics, and a couple of system-level metrics.

- **Utility:** In our context of physical activity monitoring, the first considered utility metric is the accuracy of a classifier for activity recognition. More precisely, we use the confusion matrix derived by this classifier to measure the number of correct predictions made by the classifier over all predictions made. The value of the accuracy ranges from 0 to 1, in which 1 corresponds to perfect accuracy. In addition, analytics applications monitoring physical activity usually compute and present many estimators to users. To evaluate this aspect, we compute the number of steps detected from the sanitized data and compare it with the number of steps in the raw data. To realize this, we first normalize the raw and sanitized data to compare them in the same range of values, and then compute a Peak Acceleration Threshold [1] from the raw data to estimate the number of peaks. More precisely, we used Adaptiv: An Adaptive Jerk Pace Buffer Step Detection Algorithm (<https://github.com/danielmurray/adaptiv>) for estimating the number of steps detected by the analytics application from the received data.
- **Privacy:** To assess the level of privacy of DySAN, we rely on the accuracy of inferring the sensitive attribute. In our case, we consider the gender of the users as the sensitive attribute (personal information available in public datasets). An accuracy of 0.5 corresponds to a random guess as our dataset is balanced.
- **System-level:** To assess the overhead of operating DySAN on a smartphone, we measure both the CPU time spent to sanitize the raw data on the smartphone and the energy consumption over time during a real-time processing of DySAN.

4.4 Methodology

DySAN is trained only with the MotionSense dataset while the results reported for MobiAct evaluate the transfer learning (i.e., using sanitizing models trained on MotionSense to sanitize data from MobiAct). In the training phase, we explore a range of values between 0.1 and 0.9 with a 0.1 step for both α and λ , which corresponds to 36 different sanitizing models. The sanitizer models of DySAN are trained for 300 epochs and the size of a data batch is set to 256 samples. In the online phase, we select a privacy and utility trade-off focusing primarily on privacy (i.e., ensuring the protection of the gender at the cost of the accuracy). This trade-off is controlled by the parameters x (utility) and y (privacy) (Section 3.3) which are set respectively to 0.1 and 0.9.

The random forest classifier applied during the online phase of DySAN uses a feature vector extracted from the raw signal. The choice of these descriptors was made on the basis of an earlier

²https://github.com/mmalekzadeh/motion-sense/tree/master/codes/gen_paper_codes

review on effective descriptors for gait recognition [32]. We use 4-fold cross-validation in which the testing set is randomly partitioned into 4 equal sized subsamples. Reported results correspond to average over 10 repetitions of each experiment. The computation of the different global models (each corresponding to a precise set of hyperparameters) has been parallelized on a hybrid GPU/CPU computing farm.

5 EVALUATION

In this section, we report the results obtained for the evaluation of DySAN by highlighting important features, namely the good utility and privacy trade-off (Section 5.1), the low distortion of the sanitized data (Section 5.2), the better performances compared to state-of-the-art approaches (Section 5.3), the advantage of dynamically select the best sanitizing model according to the incoming data (Section 5.4), and the limited cost of operating DySAN on a mobile (Section 5.5).

5.1 Utility and privacy trade-off

In this section, we evaluate the capacity of an analytics application to infer the gender of the user and its activity from the sanitized data provided by DySAN and sent by the mobile application. We compare the performance of several classifiers that could be used by the analytics application, namely a gradient boosting classifier (GB), a multi-layer perceptron (MLP), a long short-term memory neural network (LSTM), a decision tree (DT), a random forest (RF), a logistic regression (LR) and also two CNNs with the same architectures than the predictor and the discriminator of DySAN (referred as Raw on the figures).

Figure 3 reports the accuracy for both datasets for predicting the gender and the activity with the different classifiers as well when using the raw data. First, the results show that without any protection (*i.e.*, on raw data) the application is able to infer the gender with 98.5% accuracy. In addition, the activity is also inferred from the raw data with 97% of accuracy on average. Secondly, we can observe that DySAN successfully decreases the privacy risk with respect to inferring the sensitive attribute while limiting the loss of activity detection. Indeed, with the sanitized data, an analytics application is only able to infer the gender up to 61% and 57% of accuracy, respectively for MotionSense and MobiAct. In term of utility, depending on the classifier, the accuracy of the activity recognition varies between 78% and 92%, which represents only a small drop compared to using the use of the raw data. Remark that the LSTM, a recurrent neural network architecture commonly used for temporal signal, does not provide best results as one could expect.

5.2 Distortion of the sanitized signal

The utility of the sanitized data is not just about the activity recognition but also with respect to more fine-grained information related to the activity. In this section, we demonstrate that DySAN keeps relevant information in the signal enabling to conduct further analysis. More precisely, we consider the computation of the number of steps from the signal for MotionSense dataset. Following the step detection method presented in 4.3, Table 1 shows that with DySAN the estimation of the number of steps only suffers from a 7% error compared to the raw data. With the different baselines, the sanitized

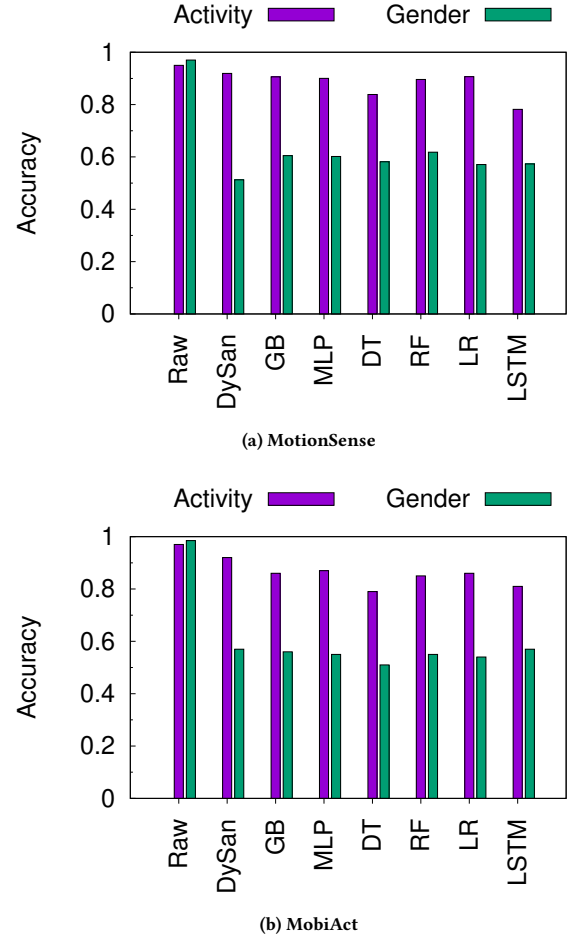


Figure 3: The sanitized data provided by DySAN drastically decreases the privacy risk compared to using the raw data while limiting the loss of activity detection, and this regardless of the classifier used.

signal appears to be much more noisy and the step detection is greatly impacted with an overestimate number of the steps of more than 64% for Olympus, more than 29 % for MSDA and more than 12% of errors for GEN. While GEN does not take into account data distortion compared to MSDA, its signal appears to be less noisy with a smaller error rate. This smaller utility loss is balanced by a smaller privacy guarantee as shown in Figure 4a. Moreover the data distortion improvement of DySAN compared to MSDA (while ensuring a better privacy) is provided by the dynamic model selection that adapts the sanitization according to the incoming signal at the user level in order to have the best utility-privacy trade-off. This fine-grain sanitization is not possible through a single sanitizing model for all users and whatever the performed activity (*i.e.*, the incoming signal). The method ORF is not considered here because it only extracts features and the signal is not preserved, which prohibits possibility to conduct further analysis.

	Steps	DTW
Raw data	14387	-
DySAN	15321 (+6.49 %)	12.96
GEN	12817 (-12.25%)	14.28
Olympus	23658 (+64.44%)	156.03
MSDA	18624 (+29.45%)	23.37

Table 1: The sanitized signal provided by DySAN appears to be less distorted and more useful for step detection than other approaches.

To evaluate the deformation of the signal, we also report the Dynamic Time Warping (DTW) [5] between the raw and the sanitized data from each baseline (Table 1). This metric measures the distortion between two temporal signals. If this metric has a small value then it means that the two signals are quite similar to each other, which is a sign of a small distortion. The results obtained show that the sanitized data produced by DySAN is more similar to the raw data compared to other baselines. Similarly to step detection, the sanitizing process of Olympus depicts a large data distortion making further analysis of the signal impossible. Other metrics assessing the deformation of the signal (*i.e.*, mean, standard deviation, skewness, kurtosis and energy) are reported in Appendix B.

5.3 Comparative analysis

We compare DySAN against baseline approaches (Figure 4). Two versions of DySAN are given to represent, DySAN where the annotations of the activities are known and the online version, DySAN(o), where the activities are not given but inferred from the random forest (RF) classifier. The first version has been added for a more fair comparison to state-of-the-art that does not evaluate models as we suggest.

For MotionSense (Figure 4a), the privacy improvement of DySAN occurs at the cost of a slightly decrease of utility (gender inference limited to 51% and an activity recognition of 92%). For the online version, which works blindly without annotations, the performance is a little worse, with a gender inference of 57% and accuracy in activity of 75%. This utility mitigation comes from the imperfect accuracy of the random forest classifier used in the online phase to select the best sanitized model. Indeed, to dynamically select the sanitizer model, DySAN needs to estimate the model providing the best utility and privacy trade-off with respect to the considered parameters (Section 3.3). To achieve this, DySAN relies on a calibration process to build a RF classifier on the smartphone using the raw data used as a reference to predict the current activity performed by the user. This local RF classifier provides an average accuracy of respectively 96% and 94% on the activity recognition for MotionSense and MobiAct datasets. While these accuracies are high, an activity wrongly predicted by this classifier leads to a selection of the sanitizer model that does not correspond to the best utility and privacy trade-off.

As depicted in Figure 4b, results for MobiAct show that DySAN and DySAN(o) outperform other approaches by limiting the gender inference to 55% and 54% while only reducing the accuracy of activity recognition by respectively 2% and 5% compared to using the raw data. Although GEN and ORF also limit significantly the

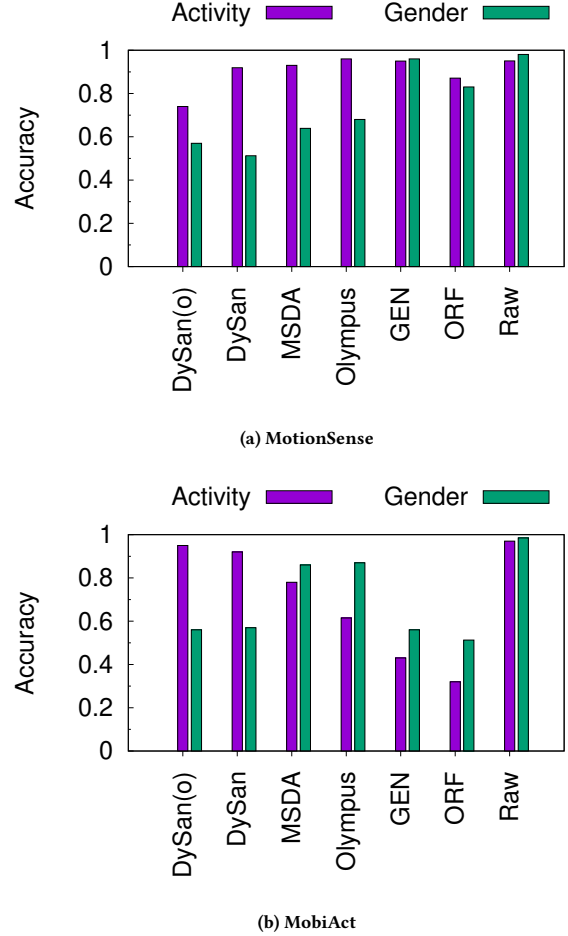


Figure 4: DySAN provides the best privacy protection compared to state-of-the-art approaches at the cost of a slightly smaller accuracy in term of activity detection.

gender inference, the accuracy of the activity detection is drastically impacted (respectively, 43% and 32%). For MSDA, both utility and privacy are not impacted compared to Olympus and GEN due to the distortion limitation that force the model to avoid large transformation on the data.

The accuracy of the classification is not uniform for all activities. Table 2 details the True Positives and False Positives of this classification for DySAN on MotionSense dataset. This table also reports the percentage of data in the dataset for each activity. We observe that the accuracy of the classification depends on the performed activity. This heterogeneity is a direct result of the unbalanced classes. Specifically, the walking activity has the highest precision which corresponds to the activity with the largest amount of data, while other activities contained less data and depicted lower good predictions. This difference in terms of good prediction between walking and other activities can also be explained by a calibration of the size window adapted for the walk (Section 2).

	TP	FP	Precision	Data percentage
Downstairs	221	112	66.4	17.2
Upstairs	223	198	53.0	20.5
Walking	918	74	92.5	44.9
Jogging	216	212	50.5	17.4

Table 2: True Positive, False Positive, Precision and percentage of data for each activity of DySan (MotionSense dataset).

Results also show the performance improvement provided by each baseline approach based on adversarial networks. Specifically, GEN, Olympus and MSDA gradually improve the utility and privacy trade-off. However, our utility analysis (Table 1) shows that the sanitized data is very distorted, which harms the possibility to perform signal processing for further analysis. MSDA integrates the data distortion in its loss function, which leads to less distorted data. This feature improves the quality of signal processing but does not significantly improve the trade-off between utility and privacy compared to Olympus (Figure 4). By dynamically selecting the best sanitizer model for each window of raw data, DySan(o) makes the gender inference close to a random guess while preserving an accurate activity detection.

The results of GEN reported in [18] mention an accuracy of 94% for the activity recognition and 64% for the gender inference for MotionSense dataset compared to 95% and 96%, respectively in our experiments. This difference comes from our implementation that exploits two neural networks for each classification task (*i.e.*, for activity recognition and gender inference) versus only one neural network for both classification tasks in the original baseline as explained in Section 4.2). However, this difference also tends to assume an over adaptation of the underlying neural network to the considered dataset. This over-adaptation is also pointed by the completely different trend for the accuracy provided for MobiAct compared to MotionSense.

5.4 Dynamic selection of sanitizing model

During the training phase, DySan computes the sanitizer models corresponding to all possible utility and privacy trade-offs by exploring the range of values for the hyperparameters α and λ . We evaluate here the benefit to dynamically adapt the sanitizing model according to the incoming data of each user compared to two static baseline approaches. Firstly, we compute the accuracy for both the gender inference and the activity recognition when the sanitizer model is fixed for all the users. This case represents the behaviors of all comparative baselines where the considered model is the one providing the best performance (*i.e.*, the utility and privacy trade-off) on average for all the users. Secondly, we consider a personalized solution in which the sanitizer model is personalized for each user. In this case, the sanitizing model is the one which provides the smallest accuracy in term of gender inference and the best accuracy in term of activity recognition according to the whole models set for a specific user. This solution provides a sanitizer model personalization but the selected model is static and does not change according to the evolution of the incoming data (and the associated changes in term of performed activity).

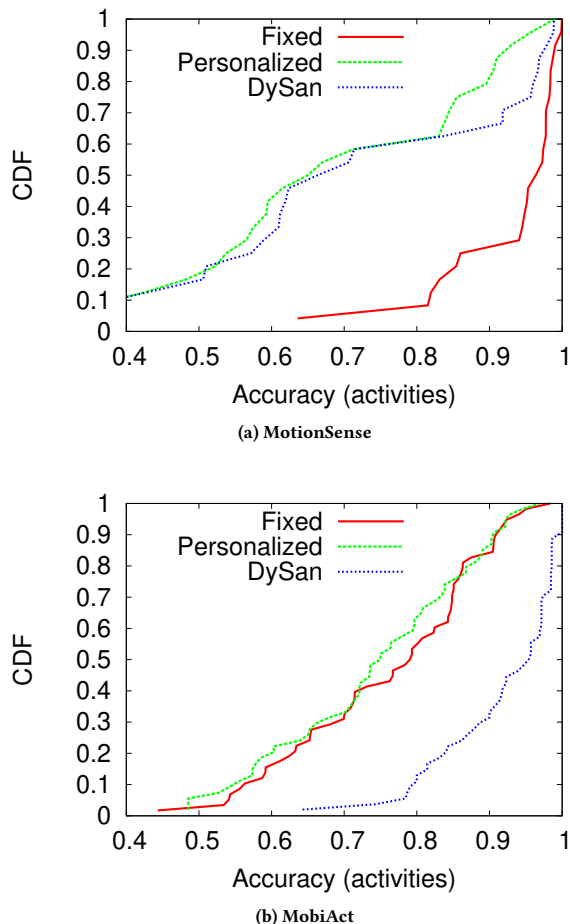


Figure 5: The dynamic sanitizing model selection of DySan significantly improves the activity recognition in case of transfer learning (*i.e.*, MobiAct dataset).

We compare these static solutions against DySan in which the considered sanitizer model for each user changes according to the incoming data in order to maximize the utility and privacy trade-off over time. Figures 5 and 6 depict for both datasets the cumulative distribution (*i.e.*, CDF) of the accuracy of the activity recognition and the gender inference respectively, when a fixed, a personalized and a dynamic sanitizing model is considered. Results show that the accuracy in both classification tasks is highly heterogeneous over the population of users. This high heterogeneity reflects the fact that a static model is not well adapted for all users or for all activities performed by the user, thus motivating the need for a dynamic approach.

Specifically, results show that dynamically adapting the sanitizing model significantly improves the activity recognition compared to using a static model in case of transfer learning (*i.e.*, MobiAct dataset, Figure 5b). For MotionSense dataset (Figure 5a), most users benefit from an important accuracy with a static model fixed for all users. This result can be explained by the fact that the sanitizing

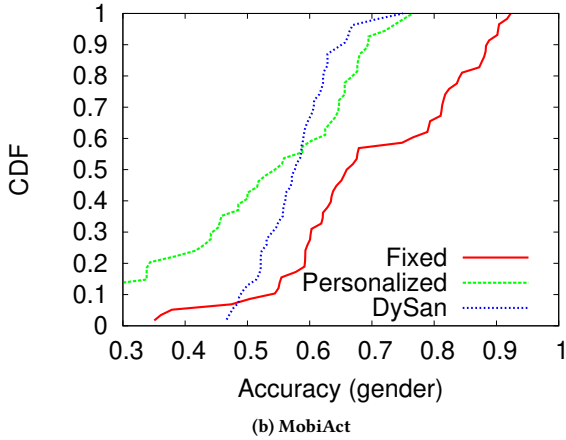
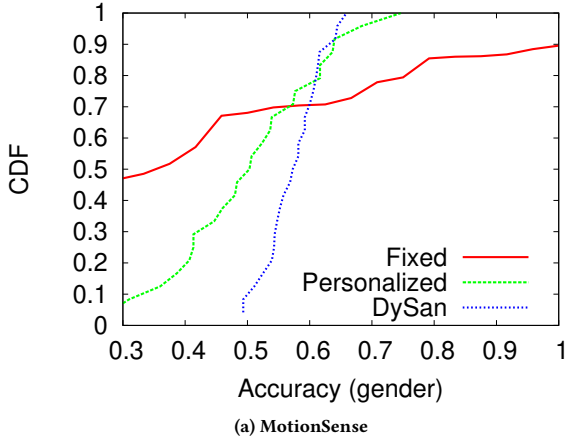


Figure 6: By dynamically adapting the sanitizing model for each user according to the incoming data, DySAN greatly improved the protection against gender inference (the distribution of the gender accuracy is more centered around 0.5, which corresponds to a random guess).

models have been learned with the same users, leading to a learning of the motion characteristics of all the considered users.

For the gender inference, the objective of the sanitizer is to provide an accuracy around 0.5 which corresponds to a random guess for all users. However, results depicted in Figure 6 clearly shows that a fixed model for all users fails to protect against gender inference. Indeed, the distribution reports a wide range of accuracy over the users where it is possible to infer the gender with 80% of confidence for 60% and 20% of the users for MobiAct and MotionSense dataset, respectively. Adopting a personalized sanitizer model for each user decreases the accuracy of the gender prediction compared to a fixed model for all users but the distribution of the accuracy is still large (from 0.3 to 0.75 for MotionSense and from 0.3 to 0.8 for MobiAct). By dynamically adapting the sanitizing model according to the incoming data, DySAN greatly improves

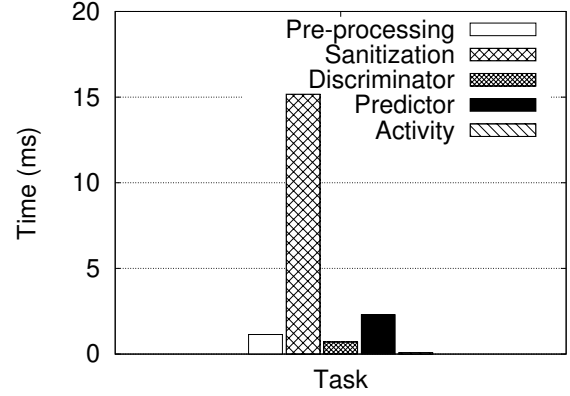


Figure 7: The limited cpu overhead of the sanitation of DySAN is compatible to real-time processing on smartphone.

the protection against gender inference compared to using a fixed model with a sharper distribution centered around 0.5.

These results also show the capacity of DySAN to transfer the learning performed on MotionSense to MobiAct (an activity recognition accuracy around 92% on average for a gender accuracy around 57%). For comparison, we evaluated the transfer learning of GEN using the original sanitizing model learned on MotionSense (and publicly available) to MobiAct dataset. In this case of transfer learning GEN provides an accuracy in terms of activity recognition and gender detection around 43% and 56%, respectively. This result shows the limited capacity of GEN to transfer learning from MotionSense to another dataset assuming an over adaptation of the underlying neural network and parameters to the considered dataset.

To complete this analysis, we evaluate the variation of the sanitizing model selection of DySAN compared to static approaches as well as the number of different models used by DySAN for each user in Appendix C. We also quantify the possibility to use the set of selected models as a fingerprint to identify each user in Appendix E.

5.5 Performance as measured on devices

We now evaluate the cost of operating DySAN on a smartphone. DySAN protects the sensitive attribute while ensuring an accurate activity recognition and minimal data distortion. However, applying the sanitizing at run time on the mobile introduces an overhead. We do not consider the overhead of the learning as it is a one time operation. DySAN evaluates multiple sanitizing models (*i.e.*, according to each α and λ hyperparameter explored) before selecting the one that produces the best compromise between utility and privacy. Consequently, the overhead associated with the sanitizing of raw data depends on the number of considered models.

Figure 7 describes the time (ms) spent by a Xiaomi Redmi Note 7 (equipped with a Qualcomm Snapdragon 660 and 3 GB of memory running a java application using Pytorch 1.6) on each task associated to a single sanitizing model of a window of incoming data (*i.e.*, 2.5 seconds of data). Specifically, these tasks include the pre-processing of signal, the sanitizing of raw data, the evaluation of the privacy and the utility on the sanitized data respectively by

	Activity accuracy (%)	Gender accuracy (%)
36 models	92	57
20 models	89	59
16 models	88	63
8 models	86	66

Table 3: Reducing the number of sanitizing models available for the selection decreases the accuracy in activity recognition while increasing the accuracy in gender inference.

the discriminator and the predictor, and the classification of the activity performed by the user from the raw data. Excepting the pre-processing, which is performed only once for a window of data, the other tasks have to be repeated for each explored sanitizing model. Results show that applying a sanitizing model once spends most of the time while all operations require 19 ms. Considering 20 or 36 sanitizing models increases this time to 366 ms and 658 ms, respectively. Although this processing is compatible with real-time processing (*i.e.*, data processed after each data window), the number of models deployed on the smartphone should be chosen to limit the overload. The number of considered sanitizer models deployed on the smartphone has also an impact on the storage space requirement. On average, the size of a single model is around 15 MB. Considering 36 models results in 540 MB which is not a limitation with regards to the storage capacities of current smartphones.

We also evaluate the impact of the considered number of sanitizing models. Considering less sanitizing models leads to cover less hyperparameter values and thus limiting the achievable utility and privacy trade-off. Consequently, a degradation of the accuracy for both the activity detection and the gender interference is observed. Table 3 presents the performances obtained with different number of sanitizing models available for the selection. Results show that from 36 to 20 sanitizing models, the accuracy in activity recognition decreases by only 3% and increase by 2% the gender inference.

Finally, we evaluate the impact of running DySAN on the energy consumption on the smartphone. Figure 8 reports the decrease in the battery charge over time for a baseline where no operation is performed on the smartphone, and for a real-time processing of DySAN (*i.e.*, after each window of raw data, and exploring 36 sanitizing models before to select the best one). In both cases, the screen remained on during the experiment. Results show that DySAN consumed 1% more battery after 1 hour, which stays a reasonable energy consumption.

6 RELATED WORK

With the availability of wearable and personal devices, there have been a growing research on the use of collected data for quantifying various aspects of personal life, such as the number of calories consumed, the blood pressure, etc. A growing literature concern the use of data for predicting the physical activities performed by users, let it be for either medical, insurance or various other reasons. We refer the interested reader to the surveys [21] and [27] on machine learning and deep learning techniques applied for predicting the type of activities performed. To address heterogeneity of data and the fact that a single learning model can degrade performance for

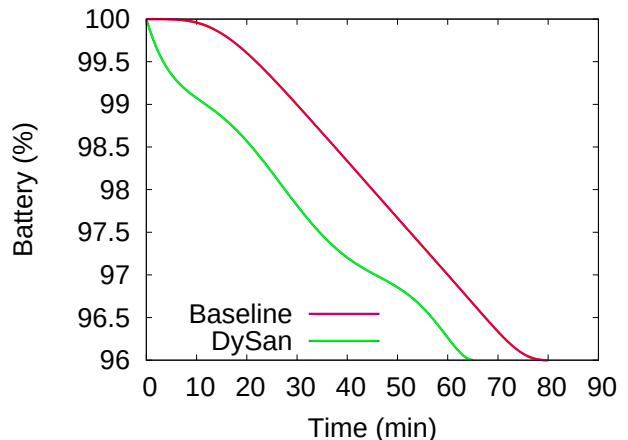


Figure 8: The impact of DySAN on energy consumption is limited (1% less battery after 1 hour).

atypical users, local adaptations of the model (*e.g.*, fine tuning) are recently considered in the context of federated learning [37].

Different approaches have been proposed to protect sensitive information in sensor data while retaining data utility. Our approach is closely related to Gansan [2], however, our framework goes beyond, by considering the data utility with respect to the Predictor network, in addition to the application on motion sensor data. Next, we focus on approaches used as baselines previously in the paper. [18] is the only one which focus on the gender as the sensitive information. [14] and [19] in their case, focus on the re-identification only while [28] apply their approach on several applications like object recognition or action recognition with several data types such as images or motion sensors. In the case adversarial approach that use autoencoders, the sensitive information can be extracted from the representation produced by the encoder [17], the decoder [28] which also correspond to our approach, or both the encoder and the decoder [19] for data sanitization. Specific to the sensor data generation, SenseGen [3] is a deep learning architecture for protecting users privacy by generating synthetic sensor data. Unfortunately, this method does not provide any guarantee on the protection.

To enlarge with other applications protecting sensitive informations using adversarial methods, [7] use a VGAN to transform face images in order to hide facial expression of the users that can be used to reveal their identity while preserving generic expressions. Adversarial approaches can also be used to hide sensitive information such as text in images [9] or identity information in the fingerprints [22].

From a broader privacy perspective, [34] proposes an adversarial network technique to minimize the amount of mutual information between a sensitive attribute and useful data while bounding the amount of distortion introduced. They applied their solution on a synthetic and a computer vision dataset. Inspired from [34], authors in [30] have developed a method for learning an optimal privacy protection mechanism also inspired from GAN, which they have

applied to location privacy. In [25], authors have proposed an approach called table-GAN, which aim at preserving privacy by generating synthetic data. By suppressing *one-to-one* relationship and limiting the quality of dataset reconstruction re-identification attacks are rendered less performant. They compared their approach with standard privacy techniques such as k -anonymity and t -closeness.

Apart from techniques using adversarial approach to protect sensitive information on sensor data, [20] proposes two privacy preserving mechanisms based on clustering algorithms called Hierarchical Agglomerative Clustering to compress amount of disclosed data so that the amount of sensitive information can be reduced. [38] in their case, develop a framework for images data made on wearable cameras that can protect sensitive information such as face, objects or locations thanks to a neural network that detects the sensitive objects which will then be blurred or deleted. Rather than focusing on re-identification, [6] investigate what data to share, in such a way that certain kinds of inferences cannot be down. They propose *ipShield* that obfuscate data according to the quantification of an adversary's knowledge regarding a sensitive inference.

7 CONCLUSION

We presented DySAN, a privacy-preserving framework which sanitizes motion sensor data in order to prevent unwanted inference of sensitive information from an adversary exploiting raw data. At the same time, DySAN preserves as much as possible the useful information for activity recognition and other estimators of physical activity monitoring. Results show that DySAN drastically reduces the risk of gender inference without impacting the ability to detect the activity or to monitor the number of steps. We also show that the dynamic sanitizing model selection of DySAN successfully adapts the protection to each user over time according to the evolution of the incoming data. This method is particularly effective in a transfer learning case unlike the other baselines that have a unique sanitizing model for all the users. Moreover, we show that the overhead introduces on the smartphone to sanitize the data is compatible with real-time processing while keeping a reasonable energy consumption. Lastly, we compared our approach with existing approaches and demonstrated that DySAN provides better control over privacy-utility trade-off. We also investigated the possibility to extend DySAN to take into account multiple sensitive attributes. Our preliminary results by adding several discriminators accounted in the loss function of the sanitizer's training are encouraging, however, we are limited by the small size of the available datasets. Indeed, making the sanitizing models more complex requires more data to capture the specificity of each use case.

ACKNOWLEDGMENTS

This work was partially funded by the ANR project PMR — Privacy-preserving methods for Medical Research (ANR-20-CE23-0013). Sébastien Gambs is supported by the Canada Research Chair program, a Discovery Grant (NSERC), the Legalia project (FQRNT) as well as a grant for the Office of the Privacy Commissioner (OPC) of Canada.

REFERENCES

[1] A. Abadleh, E. Al-Hawari, E. Alkafaween, and H. Al-Sawalqah. 2017. Step detection algorithm for accurate distance estimation using dynamic step length. In

MDM. 324–327.
 [2] Ulrich Aivodji, François Bidet, Sébastien Gambs, Rosin Claude Ngueveu, and Alain Tapp. 2019. Agnostic data debiasing through a local sanitizer learnt from an adversarial network approach. *arXiv preprint arXiv:1906.07858* (2019).
 [3] Moustafa Alzantot, Supriyo Chakraborty, and Mani Srivastava. 2017. Sensegen: A deep learning architecture for synthetic sensor data generation. In *PerCom Workshops*. 188–193.
 [4] C. BenAbdelkader, R. Cutler, and L. Davis. 2002. Stride and cadence as a biometric in automatic person identification and verification. In *IEEE FGR*. 372–377.
 [5] Donald J. Berndt and James Clifford. 1994. Using Dynamic Time Warping to Find Patterns in Time Series. In *AAAIWS*. 359–370.
 [6] Supriyo Chakraborty, Kasturi Rangan Raghavan, Matthew P. Johnson, and Mani B. Srivastava. 2013. A Framework for Context-Aware Privacy of Sensor Data on Mobile Systems. In *HotMobile*. Article 11, 6 pages.
 [7] Jiawei Chen, Janusz Konrad, and Prakash Ishwar. 2018. VGAN-Based Image Representation Learning for Privacy-Preserving Facial Expression Recognition.
 [8] Yves-Alexandre de Montjoye, César A Hidalgo, Michel Verleyesen, and Vincent D Blondel. 2013. Unique in the Crowd: The privacy bounds of human mobility. *Nature* 3 (2013).
 [9] Harrison Edwards and Amos Storkey. 2015. Censoring Representations with an Adversary.
 [10] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *SIGKDD*. 259–268.
 [11] George D. Fulk and Edward Sazonov. 2011. Using Sensors to Measure Activity in People with Stroke. *Topics in Stroke Rehabilitation* 18, 6 (2011), 746–757.
 [12] Jun Han, Emmanuel Owusu, Le T Nguyen, Adrian Perrig, and Joy Zhang. 2012. Accomplice: Location inference using accelerometers on smartphones. In *COMSNETS*. 1–9.
 [13] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery* 33, 4 (Mar 2019), 917–963.
 [14] Théo Jourdan, Antoine Boutet, and Carole Frindel. 2018. Toward privacy in IoT mobile devices for activity recognition. In *MobiQuitous*. 155–165.
 [15] Jacob Leon Kröger, Philip Raschke, and Towhidur Rahman Bhuiyan. 2019. Privacy implications of accelerometer data: a review of possible inferences. In *ICCS*. 81–87.
 [16] Seon-Woo Lee and Kenji Mase. 2002. Activity and location recognition using wearable sensors. *IEEE pervasive computing* 1, 3 (2002), 24–32.
 [17] Changchang Liu, Supriyo Chakraborty, and Prateek Mittal. 2017. DEEProtect: Enabling Inference-based Access Control on Mobile Sensing Applications.
 [18] Mohammad Malekzadeh, Richard G. Clegg, Andrea Cavallaro, and Hamed Hadadi. 2018. Protecting Sensory Data Against Sensitive Inferences. In *W-P2DS'18*. ACM, 2:1–2:6.
 [19] Mohammad Malekzadeh, Richard G. Clegg, Andrea Cavallaro, and Hamed Hadadi. 2019. Mobile sensor data anonymization. *IoTDI* (2019).
 [20] Soumia Menasria, J. Wang, and M. Lu. 2018. The purpose driven privacy preservation for accelerometer-based activity recognition. *World Wide Web* 21 (2018), 1773–1785.
 [21] Henry Friday Nweke, Ying Wah Teh, Mohammed Ali Al-Garadi, and Uzoma Rita Alo. 2018. Deep learning algorithms for human activity recognition using mobile and wearable sensor networks: State of the art and research challenges. *Expert Systems with Applications* 105 (2018), 233–261.
 [22] Witold Oleszkiewicz, Peter Kairouz, Karol Piczak, Ram Rajagopal, and Tomasz Trzcinski. 2018. Siamese Generative Adversarial Privatizer for Biometric Data.
 [23] Zhaoqing Pan, Weijie Yu, Xiaokai Yi, Asifullah Khan, Feng Yuan, and Yuhui Zheng. 2019. Recent progress on generative adversarial networks (GANs): A survey. *IEEE Access* 7 (2019), 36322–36333.
 [24] Eunjeong Park, Hyuk-Jae Chang, and Hyo Suk Nam. 2017. Use of Machine Learning Classifiers and Sensor Data to Detect Neurological Deficit in Stroke Patients. *J Med Internet Res* 19, 4 (2017), e120.
 [25] Noseong Park, Mahmoud Mohammadi, Kshitij Gorde, Sushil Jajodia, Hongkyu Park, and Youngmin Kim. 2018. Data synthesis based on generative adversarial networks. *VLDB* 11, 10 (2018), 1071–1083.
 [26] Jun Qi, Po Yang, Dina Fan, and Zhikun Deng. 2015. A survey of physical activity monitoring and assessment using internet of things technology. In *CIT/IUCC/DASC/PICOM*. 2353–2358.
 [27] Sreenivasan Ramasamy Ramamurthy and Nirmalya Roy. 2018. Recent trends in machine learning for human activity recognition—A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 8, 4 (2018).
 [28] Nisarg Raval, Ashwin Machanavajjhala, and Jerry Pan. 2019. Olympus: Sensor Privacy through Utility Aware Obfuscation. *Proceedings on Privacy Enhancing Technologies* 2019, 1 (2019).
 [29] J. L. Reyes-Ortiz. 2015. *Smartphone-based human activity recognition*. Springer.
 [30] Marco Romanelli, Catuscia Palamidessi, and Konstantinos Chatzikokolakis. 2019. Generating Optimal Privacy-Protection Mechanisms via Machine Learning. *CoRR* abs/1904.01059 (2019).

- [31] Florian Schäfer and Anima Anandkumar. 2019. Competitive gradient descent. In *Advances in Neural Information Processing Systems*. 7623–7633.
- [32] S. Sprager and M. B. Juric. 2015. Inertial sensor-based gait recognition: a review. *Sensors* 15, 9 (2015), 22089–22127.
- [33] Salvatore Tedesco, John Barton, and Brendan O’Flynn. 2017. A review of activity trackers for senior citizens: Research perspectives, commercial landscape and the role of the insurance industry. *Sensors* 17, 6 (2017), 1277.
- [34] Ardhendu Tripathy, Ye Wang, and Prakash Ishwar. 2019. Privacy-preserving adversarial networks. In *Allerton*. IEEE, 495–505.
- [35] George Vavoulas, Charikleia Chatzaki, Thodoris Malliotakis, Matthew Pedititis, and Manolis Tsiknakis. 2016. The MobiAct Dataset: Recognition of Activities of Daily Living using Smartphones. In *ICT4AgeingWell*. 143–151.
- [36] Heetae Yang, Jieun Yu, Hangjung Zo, and Munkee Choi. 2016. User acceptance of wearable devices: An extended perspective of perceived value. *Telematics and Informatics* 33, 2 (2016), 256–269.
- [37] Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. 2020. Salvaging Federated Learning by Local Adaptation.
- [38] E. Zarepour, M. Hosseini, S. S. Kanhere, and A. Sowmya. 2016. A context-based privacy preserving framework for wearable visual lifeloggers. In *PerCom*. 1–4.

APPENDICES

A NEURAL NETWORK ARCHITECTURE

In this section, we detail the underlying neural networks of DySAN³.

A.1 Discriminator Net

- (1) Input (125,6)
- (2) Conv1D (64, kernel_size=6, stride=1, activation=ReLU)
- (3) AvgPool1D(kernel_size=2, stride=2)
- (4) BatchNorm1D(100, eps=1e-05, momentum=0.1)
- (5) Dropout(p=0.5)
- (6) Dense(64, activation=ReLU)
- (7) Dense(2, activation=softmax)

A.2 Predictor Net

- (1) Input (125,6)
- (2) Conv1D (100, kernel_size=6, stride=1, activation=ReLU)
- (3) AvgPool1D(kernel_size=2, stride=2)
- (4) BatchNorm1D(100, eps=1e-05, momentum=0.1)
- (5) Conv1D(100, kernel_size=5, stride=1, activation=ReLU)
- (6) AvgPool1D(kernel_size=2, stride=2)
- (7) Conv1D(160, kernel_size=5, stride=1, activation=ReLU)
- (8) AvgPool1D(kernel_size=2, stride=2)
- (9) Conv1D(160, kernel_size=5, stride=1, activation=ReLU)
- (10) AvgPool1D(kernel_size=2, stride=2)
- (11) Dropout(p=0.5)
- (12) Dense(64, activation=ReLU)
- (13) Dense(4, activation=softmax)

A.3 Sanitizer Net

- (1) Input (125,6)
- (2) Conv1D (64, kernel_size=6, stride=1,)
- (3) Conv1D (128, kernel_size=5, stride=1)
- (4) Dense(128)
- (5) Dense(64, activation=LeakyReLU(0.01))
- (6) Dense(64)
- (7) Dense(128)
- (8) Deconv1D (128, kernel_size=5, stride=1)
- (9) Deconv1D (64, kernel_size=5, stride=1, activation=softmax)

B SANITIZED DATA DISTORTION

Table 9 gives complementary results concerning the similarity analysis of the data sanitized between the different baselines, with simple quantitative measures. Here the raw measures plus the percentage relative error are given for each baselines. Even if those metrics gives few information about the shapes of the signals, we can still observe that Olympus, the only baselines that does not

³Standard state of the art architectures (e.g., Network in Network (NiN), AlexNet, VGGNet) are not adapted for our adversarial training.

take into account the distortion of the data during training, is the one that have his measures very far from the raw data. For example the standard deviation is almost five times higher than the original data showing a very noisy signal.

C DYNAMIC SANITIZING MODE SELECTION

We evaluate the variation of the sanitizer model selection of DySAN compared to static approaches using either one model fixed for all users or one personalized model for each user. To achieved that, we measure the distance between the hyperparameters α and λ corresponding to the best privacy and utility trade-off on average for all users (i.e., the model fixed for all users) and the model selected for each user (i.e., a personalized model) or according to the incoming data (i.e., the model dynamically selected by DySAN). Figure 10 reports the distribution of this distance for both datasets. Results show that almost 40% of the users of MotionSense dataset have a personalized sanitized model which corresponds to the model providing the best trade-off on average for all users. In addition, for both datasets, results show a large variability in term of distance over all users highlighting the necessity to provide a variety of models to adapt the sanitization.

To complete this analysis, we also counted the number of different models used by DySAN for each user. Figure 11 depicted for both datasets the distribution of the percentage of all possible sanitized models (36 in our experiment as presented Section 4.4) selected by DySAN for each user. Results show a large range of number of different models selected ranging from 20% to 50%. This result show that DySAN successfully adapts the sanitization according to the evolution of the incoming data.

D UTILITY AND PRIVACY TRADE-OFF SELECTION FOR DYSAN

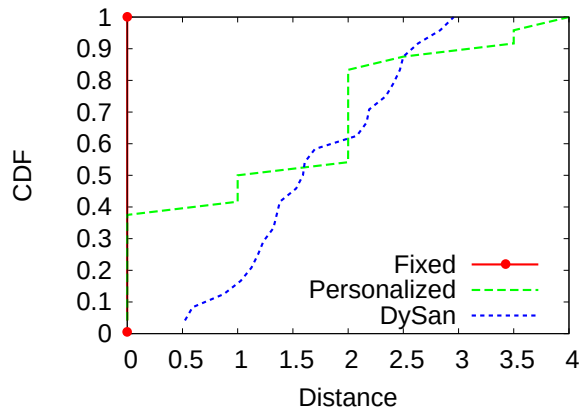
As described in Section 3.3, the best sanitizer model is selected according to the definition of the utility and privacy trade-off defined by weight coefficients x and y . Figure 12 depicts the evolution of the utility and privacy trade-off according to x and y for both datasets.

E INFORMATION LEAKAGE IN MODEL SELECTION

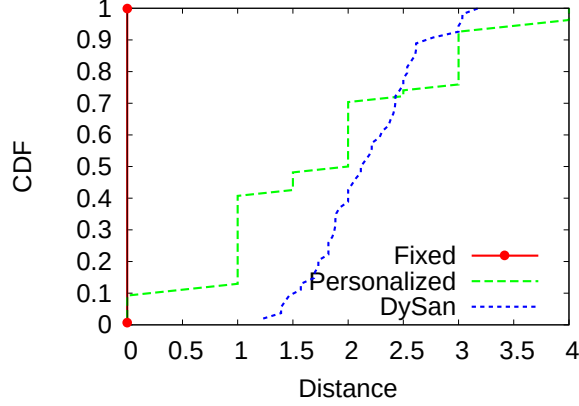
As DySAN dynamically selects the sanitizing model to use for each window of incoming data, the set of selected models could be leveraged to identify each user. Indeed, this set of sanitizing models chosen by a user could act as a unique fingerprint. To evaluate this potential information leakage, we quantify the uniqueness following the methodology presented in [8]. More precisely, the uniqueness for each user is estimated as the percentage of 100 random sets of p selected sanitizing models that are unique. Figure 13 reports for MobiAct dataset the distribution of the uniqueness with p (i.e., the size of fingerprint) from 1 to 5 and with different number of sanitizing models available for the selection. As expected, results show that the larger the fingerprint, the more unique the behaviour of a user becomes. However, at least 5 models are needed to have a strong confidence (around 80% of uniqueness) when 36 sanitizing models are exploited. To reduce this uniqueness, a lower number of sanitizing models (i.e., through the hyperparameters values explored in the training phase) should be proposed. Indeed, less

	Mean	Std	Skewness	Kurtosis	Energy
Raw	0.81	0.47	1.65	4.81	139.06
DySan	0.68 (-15.9%)	0.77 (+62.9%)	0.40 (-75.7%)	1.28 (-73.5%)	230.87 (+66.0%)
GEN	0.28 (-65.4%)	0.12 (-74.7%)	0.51 (-69.2%)	0.08 (-98.3%)	12.11 (-91.3%)
Olympus	5.40 (+566.4%)	2.52 (+433.1%)	0.61 (-62.8%)	0.29 (-94.0%)	4631.47 (+3230.5%)
MSDA	0.54 (-33.5%)	0.24 (-49.9%)	0.41 (-75.2%)	-0.11 (-102.2%)	51.87 (-62.7%)

Figure 9: Similarities metric between the raw data and the different baselines. Mean, standard deviation (std), skewness, kurtosis, energy are given in percentage of relative error.



(a) MotionSense



(b) MobiAct

Figure 10: DySan provides a large variability in terms of distance over all users highlighting the necessity to provide a variety of models to adapt the sanitization.

choice for model selection leads to have more users who share common models. Results show that exploiting less available sanitizing models reduces the uniqueness.

Information leakage in model selection leading to user re-identification is only possible if the adversary is able to characterize each selected sanitizing model from the sanitized data. In this case, the adversary could maintain a fingerprint per user to conduct its re-identification attack. To evaluate this capability, we measure

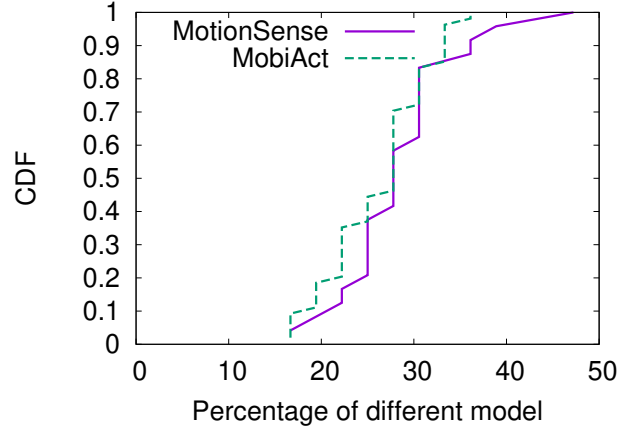


Figure 11: The data of each user is sanitized with a wide variety of models (from 20% to 50% of all the models) showing that DySan successfully adapts the sanitization according to the evolution of the incoming data.

the level of distortion using the Dynamic Time Warping of the sanitized data for each sanitizing model. Over all sanitizing models, our results show a very low standard deviation of the DTW. This low value indicates a small difference in terms of distortion when different sanitizing models are exploited, thus making it difficult for an adversary to identify the selected model from the sanitized data. This re-identification attack consequently seems difficult to achieve.

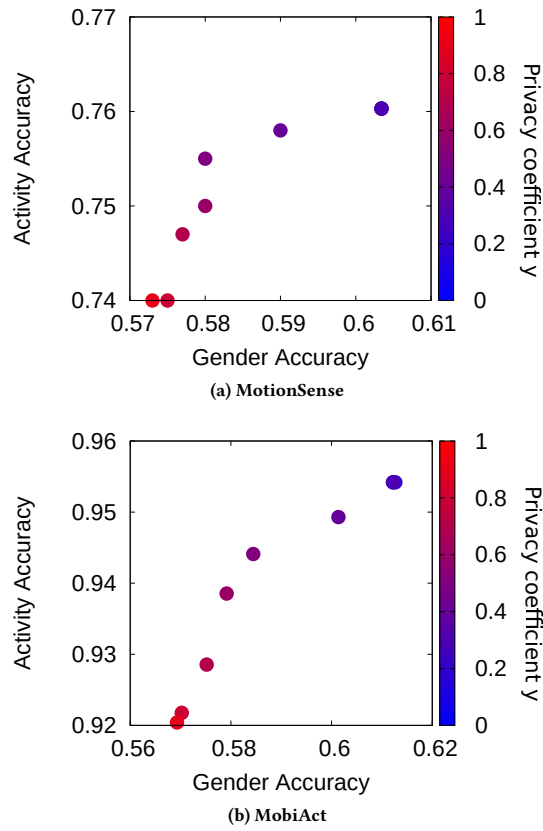


Figure 12: The variation of the Privacy coefficient y from 0.1 to 0.9 implies a variation of the trade-off between Utility and Privacy. For both dataset, when y increase, the Privacy increase (the gender accuracy decrease) and the Utility decrease (the activity accuracy decrease)

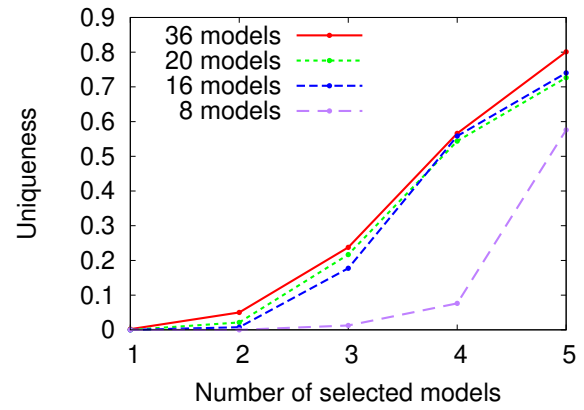


Figure 13: The uniqueness of the selected models remains low for fingerprints with less than 5 models, and depends on the number of available sanitizing models for the selection.