



HAL
open science

Deep Directed Information-Based Learning for Privacy-Preserving Smart Meter Data Release

Mohammadhadi Shateri, Francisco Messina, Pablo Piantanida, Fabrice Labeau

► **To cite this version:**

Mohammadhadi Shateri, Francisco Messina, Pablo Piantanida, Fabrice Labeau. Deep Directed Information-Based Learning for Privacy-Preserving Smart Meter Data Release. 2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (Smart-GridComm), Oct 2019, Pékin, China. pp.1-7, 10.1109/smartgridcomm.2019.8909813 . hal-04137008

HAL Id: hal-04137008

<https://hal.science/hal-04137008>

Submitted on 22 Jun 2023

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.

Deep Directed Information-Based Learning for Privacy-Preserving Smart Meter Data Release

Mohammadhadi Shateri¹, Francisco Messina¹, Pablo Piantanida^{2,3}, and Fabrice Labeau¹

¹Department of Electrical and Computer Engineering, McGill University, QC, Canada,
Email: {mohammadhadi.shateri, francisco.messina}@mail.mcgill.ca

²Laboratoire des Signaux et Systèmes, CentraleSupélec-CNRS-Université Paris Sud, Gif-sur-Yvette, France

³Montreal Institute for Learning Algorithms (Mila), Université de Montréal, QC, Canada

Abstract—The explosion of data collection has raised serious privacy concerns in users due to the possibility that sharing data may also reveal sensitive information. The main goal of a privacy-preserving mechanism is to prevent a malicious third party from inferring sensitive information while keeping the shared data useful. In this paper, we study this problem in the context of time series data and smart meters (SMs) power consumption measurements in particular. Although Mutual Information (MI) between private and released variables has been used as a common information-theoretic privacy measure, it fails to capture the causal time dependencies present in the power consumption time series data. To overcome this limitation, we introduce the Directed Information (DI) as a more meaningful measure of privacy in the considered setting and propose a novel loss function. The optimization is then performed using an adversarial framework where two Recurrent Neural Networks (RNNs), referred to as the releaser and the adversary, are trained with opposite goals. Our empirical studies on real-world data sets from SMs measurements in the worst-case scenario where an attacker has access to all the training data set used by the releaser, validate the proposed method and show the existing trade-offs between privacy and utility.

Index Terms—Privacy-preserving mechanism, Deep learning, Adversarial training, Recurrent Neural Networks, Directed information, Smart meters data.

I. INTRODUCTION

In recent decades, there has been an explosive progress in data collection tools that can measure, analyze and disseminate users data. These advances have led to a large impact in several different fields such as electrical power systems, health care, digital banking, etc. However, users are generally unwilling to share their data due to the possibility that a third party could infer their personal private information, e.g., living habits, economic status, health history. Therefore, guaranteeing users privacy while preserving the benefits of data collection is an important challenge in modern data science [1]. In this paper, we focus on this problem when the data has a time series structure and, in particular, we consider different privacy scenarios motivated by the deployment of Smart Meters (SMs) in electrical distribution networks [2]. The SMs are devices that can register a fine-grained electricity consumption of users and communicate this information to the utility provider in almost real-time. The utility of the SMs data is diverse [3], [4]. It can be used for power quality monitoring, timely fault detection, demand response, energy theft prevention, etc.

However, widespread usage of SMs data can lead to serious leakages of consumers' private information, e.g., a malicious third party could use the data to detect the presence of residents at home as well as their personal habits [2]. This problem can have a serious impact on the deployment pace of SMs and, more broadly, in the development of smart electrical grids. Thus, it is critical to ensure that SMs data are sanitized before being released.

A very simple strategy that has been proposed in the context of SMs is to use pseudonyms rather than the real identities of users for data publishing purposes [5]. However, this approach implicitly assumes that a trusted anonymizer is available. Another simple technique suggested in the literature is downsampling of the data, where the sampling rate is reduced to a level that does not pose any privacy threat [6], [7]. Although this approach may be effective from the privacy point of view, it could also limit seriously the utility of the SMs data for some applications requiring a timely response. More sophisticated and recent approaches exploit the presence of renewable energy sources and rechargeable batteries in homes to modify the actual energy consumption of users in order to hide the sensitive information [2], [8]–[11]. Some of these works use ideas from the well-known principle of differential privacy. However, recent articles suggest that the utility loss of differential privacy may be significant in practice [12]. It should be noted that our approach to the problem, which works only with the power measurement data, does not preclude the use of methods that change the energy consumption patterns by using physical resources. In fact, it should be viewed as a complementary approach that could even be used on top of the above mentioned methods.

In an information-theoretic context, privacy is generally measured by the Mutual Information (MI) between the sensitive and release variables [2], [10], [11], [13]. Some of these studies aim to find a privacy-utility trade-off using ideas from rate-distortion theory [13], [14]. More specifically, the theoretical framework of the privacy-utility problem was proposed in [13], where a hidden Markov model for the power measurements of SMs is considered in which the distribution is assumed to be controlled by only the state of the home appliances. The privacy-utility trade-off is then found for a stationary Gaussian model of electricity load with MI between

release and private sequence of variables as a privacy measure. Besides the limitation of the Gaussian model, it is noted that the MI is not well-suited to capture the causal structure of the time series data.

In this paper, an information-theoretic cost function for privacy-preserving data release of time series is proposed. In order to take into account the time series structure and causality of the data in the privacy measure, we use the Directed Information (DI) [15] between the sensitive time series and an estimation of it. Then, a cost function is derived for the releaser mechanism based on an upper bound of the DI. To optimize and validate our cost function without imposing constraints on the data distribution, two recurrent neural networks, named as releaser and adversary networks, are employed. This approach is based on the framework of Generative Adversarial Networks (GANs) [16], [17], where two neural networks are trained simultaneously with opposite goals. We will show that by controlling the relative weight between a distortion measure and the DI privacy measure, we can control the utility-privacy trade-off of SMs power measurements. A similar approach for the privacy problem, but for different applications, was considered in [14], [18] considering the standard MI and independent and identically distributed (i.i.d.) data, where the authors use two deep feed-forward neural networks for the releaser and adversary. However, to the best of our knowledge, this is the first work to consider a DI privacy measure for time series data in the general privacy-preserving context and in SMs applications in particular.

This paper is organized as follows. In Section III, we present the theoretical formulation of the problem. Then, in Section III, a privacy-preserving data release method based on Long-Short Term Memory (LSTM) Recurrent Neural Networks (RNNs) is introduced along with the training algorithm. Results for two different applications based on SMs data are presented in Section IV. Finally, some concluding remarks and a discussion about future work are given in Section V.

Notation and conventions

A sequence of random variables (X_1, \dots, X_T) of length T is denoted by X^T , while $x^T = (x_1, x_2, \dots, x_T)$ denotes a realization of X^T and $x^{(i)T} = (x_1^{(i)}, x_2^{(i)}, \dots, x_T^{(i)})$ denotes the i^{th} sample in a minibatch used for training. Mutual information [19] between variables X and Y is represented as $I(X; Y)$ and the entropy as $H(X)$. We use $X \ominus Y \ominus Z$ to indicate that X , Y and Z form a Markov chain. The expectation of a random variable X is denoted as $\mathbb{E}[X]$.

II. PROBLEM FORMULATION AND TRAINING OBJECTIVE

A. Main definitions

Consider the private variables X^T (such as occupancy label, household identity, or acorn family type), useful variables Y^T (such as actual electricity consumption of household), and observed variables W^T (which could be a combination of private and useful variables). We assume that X_t takes values on a discrete alphabet \mathcal{X} , for $t \in \{1, \dots, T\}$. A releaser \mathcal{R}_θ (this notation is used to denote that the releaser is controlled

by its parameters θ) produces the release variables as Z_t based on the observation W^t , for each time $t \in \{1, \dots, T\}$, while an adversary \mathcal{A}_ϕ attempts to infer X_t based on Z^t by finding an approximation of $p_{X^T|Z^T}$ which we shall denote by $p_{\hat{X}^T|Z^T}$. Thus, the Markov chain $(X^t, Y^t) \ominus W^t \ominus Z^t \ominus \hat{X}^t$ holds for all $t \in \{1, \dots, T\}$. In addition, due to causality, the distribution $p_{Z^T \hat{X}^T|W^T}$ can be decomposed as follows:

$$p_{Z^T \hat{X}^T|W^T}(z^T, \hat{x}^T|w^T) = \prod_{t=1}^T p_{Z_t|W^t}(z_t|w^t) p_{\hat{X}_t|Z^t}(\hat{x}_t|z^t). \quad (1)$$

The goal of the releaser \mathcal{R}_θ is to minimize the flow of information from the sensitive variables X^T to their estimation \hat{X}^T while simultaneously keeping the distortion between the release variables Z^T and the useful variables Y^T below some given value. On the other hand, the goal of the adversary \mathcal{A}_ϕ (this notation is used to denote that the adversary is controlled by its parameters ϕ) is to estimate X^T as accurately as possible.

To take into account the causal relation between X^T and \hat{X}^T , the flow of information is quantified by the DI [15]:

$$I(X^T \rightarrow \hat{X}^T) = \sum_{t=1}^T I(X^t; \hat{X}_t | \hat{X}^{t-1}), \quad (2)$$

where $I(X^t; \hat{X}_t | \hat{X}^{t-1})$ is the conditional mutual information between X^t and \hat{X}_t conditioned on \hat{X}^{t-1} [19].

The expected distortion between Z^T and Y^T is defined as:

$$\mathcal{D}(Z^T, Y^T) \triangleq \mathbb{E}[d(Z^T, Y^T)], \quad (3)$$

where $d: \mathbb{R}^T \times \mathbb{R}^T \rightarrow \mathbb{R}$ is any distortion function (i.e., a metric on \mathbb{R}^T). In order to ensure the quality of the release we shall impose the following constraint: $\mathcal{D}(Z^T, Y^T) \leq \varepsilon$ for some given $\varepsilon \geq 0$. In this work, we will consider the normalized squared error as in [13], i.e.,

$$d(z^T, y^T) \triangleq \frac{1}{T} \sum_{t=1}^T (z_t - y_t)^2. \quad (4)$$

Nevertheless, it should be noted that other distortion measures can also be relevant for the SMs data. For instance, demand response programs usually require an accurate knowledge of peak power consumption, so a distortion function closer to the infinity norm would be more meaningful for this particular application. This brief discussion simply illustrates that the distortion function should be properly matched to the intended application of the release variables Z^T in order to preserve the characteristics of the useful variables Y^T that are considered essential. Since the goal of this paper is mainly to introduce a new privacy measure and privacy-preserving data release framework, we will not further investigate different fidelity measures.

Therefore, the problem of finding an optimal releaser subject to the aforementioned adversary and distortion constraint can be formally written as follows:

$$\begin{aligned} \min_{\theta} \quad & I(X^T \rightarrow \hat{X}^T), \\ \text{s.t.} \quad & \mathcal{D}(Z^T, Y^T) \leq \varepsilon. \end{aligned} \quad (5)$$

Note that the solution of this optimization problem is a function of $p_{\hat{X}^T|Z^T}$, the conditional distributions that represent the adversary \mathcal{A}_ϕ .

B. Novel training objective

The optimization problem (5) can be directly used to define a loss function for \mathcal{R}_θ . However, note that the cost of computing the DI term is $O(|\mathcal{X}|^T)$, where $|\mathcal{X}|$ is the size of \mathcal{X} . Thus, for the sake of tractability, DI will be replaced with the following surrogate bound:

$$\begin{aligned} I(X^T \rightarrow \hat{X}^T) &= \sum_{t=1}^T H(\hat{X}_t|\hat{X}^{t-1}) - H(\hat{X}_t|\hat{X}^{t-1}, X^t) \\ &\stackrel{(i)}{\leq} \sum_{t=1}^T H(\hat{X}_t|\hat{X}^{t-1}) - H(\hat{X}_t|\hat{X}^{t-1}, Z^t) \\ &\stackrel{(ii)}{=} \sum_{t=1}^T H(\hat{X}_t|\hat{X}^{t-1}) - H(\hat{X}_t|Z^t) \\ &\stackrel{(iii)}{\leq} T \log|\mathcal{X}| - \sum_{t=1}^T H(\hat{X}_t|Z^t), \end{aligned} \quad (6)$$

where (i) is due to the fact that conditioning reduces entropy; equality (ii) is due to the Markov chains $X^t \ominus Z^t \ominus \hat{X}^t$ and $\hat{X}^{t-1} \ominus Z^t \ominus \hat{X}^t$; and (iii) is due to the trivial bound $H(\hat{X}_t|\hat{X}^{t-1}) \leq H(\hat{X}_t) \leq \log|\mathcal{X}|$. Therefore, the loss function for \mathcal{R}_θ can be written as

$$\mathcal{L}_{\mathcal{R}}(\theta, \phi, \lambda) = \mathcal{D}(Z^T, Y^T) - \frac{\lambda}{T} \sum_{t=1}^T H(\hat{X}_t|Z^t), \quad (7)$$

where $\lambda \geq 0$ controls the privacy-utility trade-off and the factor $1/T$ has been introduced for normalization purposes. It should be noted that the value of λ in (7) indirectly controls the achievable ε in (5), which means that we can control the privacy-utility trade-off by varying λ . For $\lambda = 0$, the loss function $\mathcal{L}_{\mathcal{R}}(\theta, \phi, \lambda)$ reduces to the expected distortion, being independent from the adversary \mathcal{A}_ϕ . In such scenario, \mathcal{R}_θ offers no privacy guarantees. Conversely, for very large values of λ , the loss function $\mathcal{L}_{\mathcal{R}}(\theta, \phi, \lambda)$ is dominated by the upper bound on the DI, so that privacy is the only goal of \mathcal{R}_θ . In this regime, we expect the adversary \mathcal{A}_ϕ to completely fail in the task of estimating X^T , i.e., to approach to random guessing performance.

On the other hand, the adversary \mathcal{A}_ϕ is a classifier which optimizes the following cross-entropy loss:

$$\mathcal{L}_{\mathcal{A}}(\phi) = \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[-\log p_{\hat{X}_t|Z^t}(X_t|Z^t) \right], \quad (8)$$




Fig. 1. Privacy-Preserving framework. The seed noise U^T is generated from i.i.d. samples according to a uniform distribution: $U_t \sim U[0, 1]$.

where the expectation should be understood w.r.t. $p_{X_t|Z^t}$. Notice that

$$\frac{1}{T} \sum_{t=1}^T H(X_t|Z^t) \leq \mathcal{L}_{\mathcal{A}}(\phi). \quad (9)$$

Therefore, if the adversary is ideal (i.e., $p_{\hat{X}_t|Z^t} = p_{X_t|Z^t}$ for all t), the releaser network, by maximizing $\frac{1}{T} \sum_{t=1}^T H(\hat{X}_t|Z^t)$, prevents the adversary to infer private data.

III. PRIVACY-PRESERVING MECHANISM

Based on the previous theoretical formulation, an adversarial modeling framework consisting of two RNNs, a releaser \mathcal{R}_θ and an adversary \mathcal{A}_ϕ , is considered (see Fig. 1). Note that independent noise U^T is appended to W^T in order to randomize the released variables Z^T , which is a popular approach in privacy-preserving methods. In addition, the available theoretical results show that, for Gaussian distributions, the optimal release contains such a noise component [13], [14]. For both networks, a LSTM architecture is selected (see Fig. 2), which was shown to be successful in several problems dealing with sequences of data (see [20] and references therein for more details). The training of the suggested framework is performed using Algorithm 1 which uses k gradient steps to train \mathcal{A}_ϕ followed by one gradient step to train \mathcal{R}_θ . Note that k should be large enough to ensure that \mathcal{A}_ϕ is a strong adversary during training. This is in fact the common practice for effectively training two networks in an adversarial framework [16] and, in particular, in privacy scenarios [14]. It should be recalled that, after the training of both networks is completed, an attacker network is trained in order to test the privacy achieved by the releaser network. It should be clarified that this attacker network is distinct from the adversary network used during training and illustrated in Fig. 1 the attacker used in testing mimics a real-world attacker that would try to deduce the private data from the release data.

IV. RESULTS AND DISCUSSION

A. Description of datasets

In this study, the Electricity Consumption & Occupancy (ECO) and Pecan Street data sets are used. The ECO data set, collected and published by [21], includes 1 Hz power consumption measurements and occupancy information of five

Algorithm 1 Algorithm for training privacy-preserving data releaser neural network.

Input: Data set (which includes sample sequences of useful data y^T , sensitive data x^t); seed noise samples u^T ; seed noise dimension m ; batch size B ; number of steps to apply to the adversary k ; gradient clipping value C ; L_2 recurrent regularization parameter β .

Output: Releaser network \mathcal{R}_θ .

- 1: **for** number of training iterations **do**
 - 2: **for** k steps **do**
 - 3: Sample minibatch of B examples: $\mathcal{B} = \{w^{(b)T} = (x^{(b)T}, y^{(b)T}, u^{(b)T}); b = 1, 2, \dots, B\}$.
 - 4: Compute the gradient of $\mathcal{L}_A(\phi)$, approximated with the minibatch \mathcal{B} , w.r.t. to ϕ .
 - 5: Update the adversary by applying the RMSprop optimizer with clipping value C .
 - 6: **end for**
 - 7: Sample minibatch of B examples: $\mathcal{B} = \{w^{(b)T} = (x^{(b)T}, y^{(b)T}, u^{(b)T}); b = 1, 2, \dots, B\}$.
 - 8: Compute the gradient of $\mathcal{L}_R(\theta, \phi, \lambda)$, approximated with the minibatch \mathcal{B} , w.r.t. to θ .
 - 9: Use Ridge(L_2) recurrent regularization with value β and update the releaser by applying RMSprop optimizer with clipping value C .
 - 10: **end for**
-




Fig. 2. LSTM recurrent network cell diagram. The cell includes four gating units to control the flow of information. All the gating units have a sigmoid activation function (σ) except for the input unit (that uses a hyperbolic tangent activation function (\tanh) by default). The parameters b, V, W are respectively biases, input weights, and recurrent weights. In the LSTM architecture, the forget gate $f_t = \sigma(b^f + K^f h_{t-1} + V^f w_t)$ uses the output of the previous cell (which is called hidden state h_{t-1}) to control the cell state C_t to remove irrelevant information. On the other hand, the input gate $g_t = \sigma(b^g + K^g h_{t-1} + V^g w_t)$ and input unit adds new information to C_t from the current input. Finally, the output gate $o_t = \sigma(b^o + K^o h_{t-1} + V^o w_t)$ generates the output of the cell from the current input and cell state.

houses in Swiss over a period of 8 months. In this study we re-sampled the data to have hourly samples. On the other hand, the Pecan Street data set contains hourly SMS data of houses in Texas, Austin and was collected by Pecan Street Inc. [22]. Pecan Street project is a smart grid demonstration research program which provides electricity, water, natural gas, and solar energy generation measurements for over 1000 houses in Texas, Austin. In order to model time dependency over each day (with a data rate of 1 sample per hour), the data was reshaped to sample sequences of length 24. For the ECO and Pecan Street data set, a total number of 11225 and 9120 sample sequences are used, respectively. The data is splitted into train and test sets with a ratio of roughly 85 : 15 while 10% of the training data is used as the validation set. It should

be noted that in this study, we assume that the attacker has access to all the training data used by the releaser, which can be considered as a worst-case scenario study.

B. Inference of households occupancy

The first practical case of study regarding privacy-preserving in time series data is the concern of inferring presence/absence of residents at home from the total power consumption collected by SMS [23], [24]. For this application, the electricity consumption measurements from the ECO data set are considered as the useful data, while occupancy labels are considered as the private data. Therefore, our privacy-preserving data release method aims to minimize a trade-off between the distortion of the total electricity consumption incurred and the probability of inferring the presence of an individual at home from the release signal. The releaser and adversary networks used for the training consist of 4 LSTM layers with 64 cells and 2 LSTM layers with 32 cells, respectively where a tanh activation function is used. In addition, recurrent regularizer with parameter $\beta = 1.5$ was used in each layer of the release network. The values of the other hyperparameters (B, k, m) were set to (128, 4, 8), respectively. Finally, after training, a strong attacker is used, consisting of 3 LSTM layers.

Based on the target data Y^T and the release data Z^T , the normalized root mean-square-error (NRMSE) is defined by

$$\text{NRMSE} \triangleq \sqrt{\frac{\mathbb{E}[\|Y^T - Z^T\|^2]}{\mathbb{E}[\|Y^T\|^2]}}. \quad (10)$$

Fig. 3 shows the empirically found privacy-utility trade-off for this application. It can be seen that by adding more distortion on the released data, the attacker is pushed toward a random guessing classifier.

In order to provide more insights about the release mechanism, the Power Spectrum Density (PSD) of the input signal and the PSD of the error signal (defined as the difference

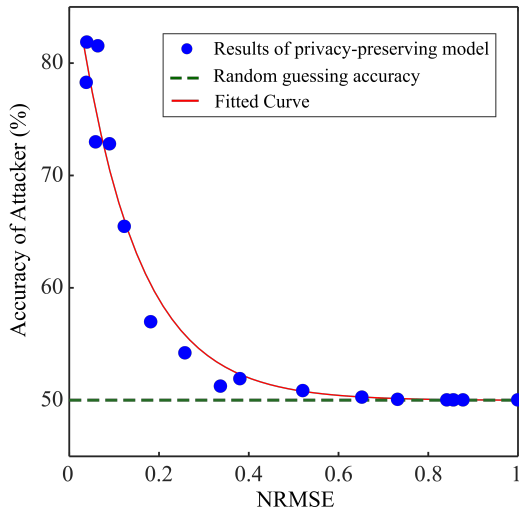


Fig. 3. Privacy-utility trade-off for house occupancy inference application. Since in this application the attacker is a binary classifier, the random guessing (balanced) accuracy is 50%. The fitted curve is based on an exponential function and is included only for illustration purposes.

between the actual power consumption and the released signal) for four different cases along the privacy-utility trade-off curve of Fig. 3 are estimated using Welch’s method [25]. For each case, we use 10 release signals and average the PSD estimates. Results are shown in Fig. 4. Looking at the PSD of the input signal (useful data) some harmonics are visible. The PSD of the error signals show that the model controls the trade-off in privacy-utility by mainly modifying the distortion on these harmonics.

It should be mentioned that two stationary tests including the Augmented Dickey-Fuller test [26] and the Kwiatkowski, Phillips, Schmidt, and Shin (KPSS) test [27] applied to our data set indicates that there is enough evidence to suggest the data is stationary, supporting our PSD analysis.

C. Inference of house identity

The second practical case of study regarding the privacy-preserving in SMs measurements is identity recognition from total power consumption of households [5]. It is assumed that the attacker has access to total power consumption of different households in a region (training data) and then attempts to determine identities of the households using the new released data (test data). Thus, our model aims at generating release data of total power consumption of households in a way that prevents the adversary to perform the identity recognition while keeping distortion on the total power minimized. For this task, total power consumption of five houses is used. For this application, the releaser consists of 6 LSTM layers each includes 128 cells and adversary has 4 LSTM layers with 32 cells. Similarly, a tanh activation function was applied and $\beta = 2$ was used in each layer of the release network. The values of the other hyperparameters (B, k, m) are set

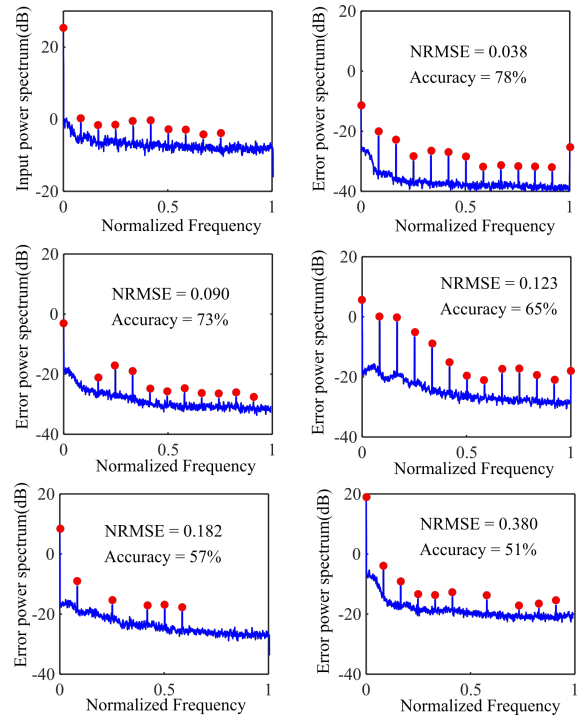


Fig. 4. PSD of the actual electricity consumption and error signals for the house occupancy inference application.

to $(128, 5, 3)$, respectively. Finally, after training, an attacker, consisting of 4 LSTM layers was used. The empirical privacy-utility trade-off curve obtained for this application is presented in Fig. 5. Comparing Fig. 5 with Fig. 3 we see that a high level of privacy is expensive. For instance, in order to obtain an attacker accuracy of 30 %, the NRMSE should be approximately equal to 0.30. This is attributed to the fact that this task is harder from the learning point of view than the one considered in Section IV-B.

PSD analysis was also performed for this application, yielding the results of Fig. 6. Once again, we see that the release network provides privacy-utility trade-off by mainly distorting the harmonics on the actual electricity consumption signal.

V. CONCLUSION

We have presented a new method to train privacy-preserving mechanisms controlling the privacy-utility trade-off in time series data. This lead us to define the directed information between sensitive variables and their estimation as a more suitable privacy measure than previous proposals in the literature. A tractable upper bound was then derived and a deep learning adversarial framework between two recurrent neural networks was introduced to optimize the new loss function. Our method was validated with two well-known privacy problems in smart meters data using two different open data sets. For both privacy problems we considered the worst-case where an attacker has access to all the training data used by the releaser. In future work, we will consider alternative formulations of the




Fig. 5. Privacy-utility trade-off for house identity inference application. Since in this application the attacker is a five-class classifier, the random guessing (balanced) accuracy is 20%. The fitted curve is based on an exponential function and is included only for illustration purposes.

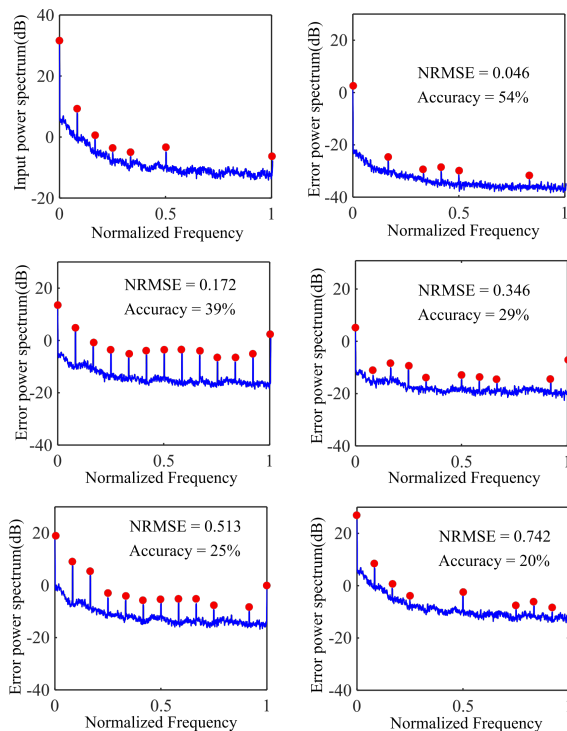


Fig. 6. PSD of the actual electricity consumption and error signals for the house identity inference application.

problem such as different distortion measures and a more general loss function in order to attempt to provide universal privacy guarantees (i.e., independent of the attacker structure and computational power).

ACKNOWLEDGMENT

This work was supported by Hydro-Quebec, the Natural Sciences and Engineering Research Council of Canada, and McGill University in the framework of the NSERC/Hydro-Quebec Industrial Research Chair in Interactive Information Infrastructure for the Power Grid (IRCPJ406021-14). This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 797805.

REFERENCES

- [1] B. Nelson and T. Olovsson, "Security and privacy for big data: A systematic literature review," in *2016 IEEE International Conference on Big Data (Big Data)*, pp. 3693–3702, Dec 2016.
- [2] G. Giaconi, D. Gunduz, and H. V. Poor, "Privacy-aware smart metering: Progress and challenges," *IEEE Signal Processing Magazine*, vol. 35, no. 6, pp. 59–78, 2018.
- [3] Y. Wang, Q. Chen, T. Hong, and C. Kang, "Review of smart meter data analytics: Applications, methodologies, and challenges," *IEEE Transactions on Smart Grid*, 2018.
- [4] S. S. R. Depuru, L. Wang, V. Devabhaktuni, and N. Gudi, "Smart meters for power grid — challenges, issues, advantages and status," in *2011 IEEE/PES Power Systems Conference and Exposition*, pp. 1–7, March 2011.
- [5] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *2010 First IEEE International Conference on Smart Grid Communications*, pp. 238–243, IEEE, 2010.
- [6] A. Cárdenas, S. Amin, and G. Schwartz, "Privacy-aware sampling for residential demand response programs," *Proceedings of 1st international ACM*, url: <http://www.eecs.berkeley.edu/schwartz/HiCons2012ASG.pdf>, 2012.
- [7] D. Mashima, "Authenticated down-sampling for privacy-preserving energy usage data sharing," in *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 605–610, IEEE, 2015.
- [8] M. Backes and S. Meiser, "Differentially private smart metering with battery recharging," in *Data Privacy Management and Autonomous Spontaneous Security*, pp. 194–212, Springer, 2013.
- [9] J. Zhao, T. Jung, Y. Wang, and X. Li, "Achieving differential privacy of data disclosure in the smart grid," in *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*, pp. 504–512, IEEE, 2014.
- [10] S. Li, A. Khisti, and A. Mahajan, "Information-theoretic privacy for smart metering systems with a rechargeable battery," *IEEE Transactions on Information Theory*, vol. 64, no. 5, pp. 3679–3695, 2018.
- [11] E. Erdemir, P. L. Dragotti, and D. Gunduz, "Privacy-cost trade-off in a smart meter system with a renewable energy source and a rechargeable battery," *arXiv preprint arXiv:1902.07739*, 2019.
- [12] R. Mendes and J. P. Vilela, "Privacy-preserving data mining: methods, metrics, and applications," *IEEE Access*, vol. 5, pp. 10562–10582, 2017.
- [13] L. Sankar, S. R. Rajagopalan, S. Mohajer, and H. V. Poor, "Smart meter privacy: A theoretical framework," *IEEE Transactions on Smart Grid*, vol. 4, pp. 837–846, June 2013.
- [14] A. Tripathy, Y. Wang, and P. Ishwar, "Privacy-preserving adversarial networks," *arXiv preprint arXiv:1712.07008*, 2017.
- [15] J. Massey, "Causality, feedback and directed information," in *Proc. Int. Symp. Inf. Theory Applic. (ISITA-90)*, pp. 303–305, Citeseer, 1990.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27 (Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds.)*, pp. 2672–2680, Curran Associates, Inc., 2014.
- [17] C. Huang, P. Kairouz, X. Chen, L. Sankar, and R. Rajagopal, "Context-aware generative adversarial privacy," *Entropy*, vol. 19, no. 12, p. 656, 2017.

- [18] C. Feutry, P. Piantanida, Y. Bengio, and P. Duhamel, "Learning Anonymized Representations with Adversarial Neural Networks," *arXiv e-prints*, p. arXiv:1802.09386, Feb 2018.
- [19] T. M. Cover and J. A. Thomas, "Elements of information theory, 2nd edition," *Wiley-Interscience: NJ*, 2006.
- [20] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>
- [21] C. Beckel, W. Kleiminger, R. Cicchetti, T. Staake, and S. Santini, "The eco data set and the performance of non-intrusive load monitoring algorithms," in *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, pp. 80–89, ACM, 2014.
- [22] Pecan Street Inc, "Dataport: the world's largest energy data resource," 2019. <https://dataport.cloud/>
- [23] W. Kleiminger, C. Beckel, and S. Santini, "Household occupancy monitoring using electricity meters," in *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*, pp. 975–986, ACM, 2015.
- [24] W. Jia, H. Zhu, Z. Cao, X. Dong, and C. Xiao, "Human-factor-aware privacy-preserving aggregation in smart grid," *IEEE Systems Journal*, vol. 8, no. 2, pp. 598–607, 2014.
- [25] P. Stoica and R. Moses, *Spectral Analysis of Signals*. Pearson Prentice Hall, 2005.
- [26] D. A. Dickey and W. A. Fuller, "Distribution of the estimators for autoregressive time series with a unit root," *Journal of the American statistical association*, vol. 74, no. 366a, pp. 427–431, 1979.
- [27] D. Kwiatkowski, P. C. Phillips, P. Schmidt, and Y. Shin, "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?," *Journal of econometrics*, vol. 54, no. 1-3, pp. 159–178, 1992.