Considering Time Correlation in the Estimation of Privacy Loss for Consumers with Smart Meters

Jun-Xing Chin*, Giulio Giaconi‡, Tomas Tinoco De Rubira†, Deniz Gündüz‡, and Gabriela Hug*
* Power Systems Laboratory, ETH Zurich, Zurich, Switzerland
† Grid Operations and Planning, Electric Power Research Institute, Palo Alto, California, USA
‡ Imperial College London, London, UK
Emails: {chin, hug}@eeh.ee.ethz.ch, {g.giaconi, d.gunduz}@imperial.ac.uk, ttinoco@epri.com

***Abstract*—Global electricity smart meter roll-out has brought about serious privacy risks for consumers. The masking of consumer consumption using rechargeable batteries has been studied as a means of protecting consumer privacy. One metric used to measure the effectiveness of such approaches is the empirical mutual information (MI), whose computation requires the estimation of both consumer load and grid-visible load distributions. These distributions have previously been modelled as independent and identically distributed (i.i.d.), or as stationary first-order Markov processes for simplicity. However, consumer load statistics are time-varying in nature, and have inherent inter-temporal dependencies. Consequently, the empirical MI based on the stationarity assumption lacks accuracy, resulting in the risk of underestimating the information leakage. In this paper, we propose using features to characterise the change in consumer demand, modelling them as feature-dependent first-order Markov processes to better approximate the actual privacy-loss. Results indicate that this approach is more accurate than i.i.d. models, and in certain cases may be a better empirical estimate of MI compared to stationary first-order Markov models.**

***Index Terms*—consumer privacy, energy management, Markov process, smart meter, time correlation**

I. Introduction

The global roll-out of electricity smart meters (SMs) has been touted as a means to enable real-time monitoring of distribution grids, allowing for more efficient grid management and planning. However, as SMs provide high-frequency consumption measurements, they also entail serious privacy risks for consumers. In fact, detailed electricity consumption load profiles reveal highly private information about consumers, such as their habits, presence at home and working hours, potential illnesses, and the equipment being used [1], [2]. A recent survey in the US has shown that utilities pose high privacy risks, and are not highly trusted by consumers [3].

Several methods have been studied in the literature to protect consumer privacy, including the manipulation of SM measurement data (e.g., aggregation and noise addition), and physically shaping the original power consumption before reporting it to the energy provider. One approach among the latter techniques masks a consumer's energy consumption, *i.e.*, the *consumer load*, by means of a rechargeable energy storage component, so that the electricity consumption reported to the energy provider, *i.e.*, the *grid load*, is different from the consumer load [4]. In [5], a heuristic battery control policy is used to hide the actual consumption by attempting to achieve a levelled load, while [6] limits the grid load to a finite number of energy levels in order to hide consumer behaviour.

However, privacy protection alone is unlikely to justify the cost of investment in energy storage devices, which according to estimates in [7] remains high. Hence, joint privacy-loss and energy cost minimisation is studied in [8] and [9]. In [9], an algorithm based on model-distribution predictive control (MDPC) is proposed to minimise privacy loss and energy cost by solving optimisation problems in a receding horizon manner. The MDPC controller uses the empirical mutual information (MI) between the consumer load and the grid load as a measure of privacy loss in its objective function, estimated by assuming that both the consumer and grid loads are stationary, independent and identically distributed (i.i.d.). Binary variables are then introduced to count the predicted observations in order to estimate this approximated MI in the MDPC controller's objective function, resulting in the controller solving mixed-integer quadratic programs whenever new meter readings are available.

While measuring privacy by this approximated MI may be adequate for obtaining privacy-protecting control actions, it is inadequate for the precise evaluation and comparison of privacy-protecting methods. In particular, one key drawback of the analysis in [9] is that MI is approximated by neglecting time correlations in the consumer and grid loads, which can lead to an underestimation of the privacy-loss, as shown later in this paper.

Accurate sample-based estimation of MI is an active area of research in the field of computer science, but recent works such as [10] and [11] require large amounts of data that is usually lacking for individual consumers. Moreover, they rely on machine learning techniques that can be less readily integrated into the objective function of control policies for use in home energy management units. In [6] and [8], an attempt is made to capture the time correlation by assuming that the consumer load, and the joint consumer and grid load statistics are stationary first-order Markov distributions. Despite being able to overcome the issue of data scarcity, the stationarity

This work was supported in part by the *Swiss National Science Foundation* and the *UK Engineering and Physical Sciences Research Council (EPSRC)* through the COPES project under the CHIST-ERA Resilient Trustworthy Cyber-Physical Systems (RTCPS) initiative.

assumption has been shown to lead to poor performance when modelling feature-dependent components of consumer load [12]. In this work, we propose expanding the stationary first-order Markov models used in [6] and [8] to account for the feature-dependent nature of the consumer load statistics by modelling them as a non-stationary first-order Markov process. We classify data into samples of different random variables, which allows us to estimate the time-varying Markov process using limited observations.

The rest of this paper is structured as follows. While Section II introduces the general problem description, Section III provides a brief introduction to the MI approximation method used in [9]. In Section IV, we introduce the proposed method based on feature-dependent first-order Markov chains. Numerical experiments are presented in Section V. Section VI concludes this paper, providing an outlook for future work.

II. Problem Formulation

Let $X_t \in \mathcal{X}_t$ and $Y_t \in \mathcal{Y}_t$ denote the consumer and grid loads at time t, respectively, where $\mathcal{X}_t \subset \mathbb{R}_+$ and $\mathcal{Y}_t \subset \mathbb{R}_+$ denote the corresponding alphabets (their domains). Consumer privacy can be measured via the MI between the consumer and grid loads [6], [8], which quantifies the amount of information shared between the two random processes. The average MI between the consumer and grid load is formulated as

$$\frac{1}{t} I(X^t;Y^t) := \frac{1}{t} \int_{X^t \in \mathcal{X}^t} \int_{Y^t \in \mathcal{Y}^t} p_{X^t,Y^t}(x^t,y^t) \times \log \frac{p_{X^t,Y^t}(x^t,y^t)}{p_{X^t}(x^t) p_{Y^t}(y^t)}\, dy^t\, dx^t, \quad (1)$$

where $X^t := (X_1,\dots,X_t)$, $Y^t := (Y_1,\dots,Y_t)$, $p_{X,Y}$, p_X and p_Y denote the probability density functions of (X,Y), X and Y, respectively, and $\log$ denotes the base-2 logarithm. For the rest of the paper, we use the shorthand $p(a)$ to denote $p_A(A=a)$ for a random variable A. The objective of a privacy-protecting policy is to minimise (1). However, it is often impossible to evaluate this function because the probability distribution functions are not readily available; hence, the need to approximate MI.

III. MI Approximation Based on I.I.D. Assumption

In [9], MI in (1) is approximated by assuming that both the consumer and grid load distributions are stationary and i.i.d. with distribution $p_{X,Y}(x,y)$ over all observed samples, and by considering a discrete time model. Furthermore, both loads are assumed to have finite support, *i.e.*, $\mathcal{X}$ and $\mathcal{Y}$ are finite. The approximate MI is then estimated by replacing probabilities with estimates based on the empirical relative frequencies, *i.e.*,

$$\begin{aligned}\frac{1}{t} I(X^t;Y^t) &\approx \frac{1}{t} \sum_{\tau=1}^{t} I(X_\tau;Y_\tau) \\ &= \sum_{X\in\mathcal{X}} \sum_{Y\in\mathcal{Y}} p(x,y) \log \frac{p(x,y)}{p(x)p(y)}, \quad (2)\end{aligned}$$

$$p(a) \approx \frac{\delta_a}{N_A}, \quad (3)$$

where δ_a is the total number of observations of $A=a$; N_A is the total number of realised samples of variable A; and (2) follows from the i.i.d. assumption, *i.e.*, $I(X_1;Y_1) = \cdots = I(X_t;Y_t) = I(X;Y)$. Note that in (2) t denotes a discrete time step. While this approximation may be sufficient to obtain a privacy-protecting control policy, it fails to capture the temporal time correlation of the consumer and grid loads when used as an evaluation metric to compare different privacy-protection methods, which is crucial when assessing the privacy-loss.

IV. Approximating MI using Feature-Dependent First-Order Markov Chains

In [8], an upper bound for (1) that captures some of the time correlation present in consumer and grid loads was proposed by also considering a discrete time model; by assuming that only the consumer load X_t and the pair $\left(X_t, Y_t\right)$, $\forall\, t \in \mathbb{Z}_+$, are stationary first-order Markov processes; and that $\mathcal{X}$ and $\mathcal{Y}$ are finite and identical for different time steps, *i.e.*,

$$p_{X^t} = p_{X_1} \prod_{\tau=2}^{t} p_{X_\tau|X_{\tau-1}},$$

$$p_{X^t,Y^t} = p_{X_1,Y_1} \prod_{\tau=2}^{t} p_{X_\tau,Y_\tau|X_{\tau-1},Y_{\tau-1}}.$$

This allows the introduction of an upper bound on $I(X^t;Y^t)$:

$$\begin{aligned}\frac{1}{t} I(X^t;Y^t) &\le \frac{1}{t} \left[\sum_{\tau=2}^{t} I(X_\tau, X_{\tau-1}; Y_\tau, Y_{\tau-1}) - \sum_{\tau=3}^{t} I(X_{\tau-1};Y_{\tau-1}) \right] \quad (4) \\ &= \frac{1}{t} \Big[(t-1)\, I(X_t, X_{t-1}; Y_t, Y_{t-1}) - (t-2)\, I(X_{t-1};Y_{t-1}) \Big]. \quad (5)\end{aligned}$$

Limitations of this approach are that it makes the assumption of stationarity, and that it uses a first-order Markov assumption. While consumer loads can be more realistically modelled as non-stationary variable-order Markov processes, it has been shown in [13] that models based on non-stationary first-order Markov processes are nonetheless sufficiently accurate for generating synthetic consumer load profiles, *i.e.*, it is not necessary to have a higher order model; nonetheless, the non-stationarity is an important characteristic that should not be neglected. However, modelling approaches based on non-stationary first-order Markov processes treat realisations at different time instances as distinct random variables; and hence, require large datasets for estimating their distributions. These large datasets are typically only obtained by grouping data from multiple consumers. For this reason, these approaches are not well suited for modelling the consumption of individual consumers, which is necessary when studying the performance of a home energy management unit that is supposed to protect the private information of the consumer. Hence, we propose reducing the number of probability distribution estimates, and thus, the required amount of data, by assuming that a consumer

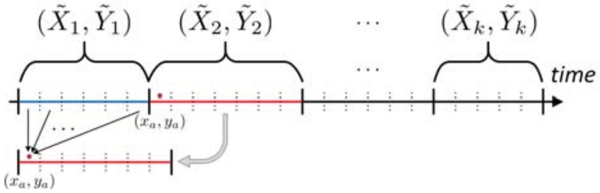
$(\tilde{X}_1, \tilde{Y}_1)$
$(\tilde{X}_2, \tilde{Y}_2)$
$\cdots$
$(\tilde{X}_k, \tilde{Y}_k)$
time
(x_a, y_a)
(x_a, y_a)

Figure 1. Transition probability estimation through counting.

can be represented by a smaller finite set of random variables, which are described by a feature set. For example, such features might be time-of-day, day-of-week, season, ambient temperature and solar irradiation.

In a time series of consumer load values, realisations having the same values of the features are assumed to be i.i.d. samples of the corresponding random variables, and we further assume that all these random variables follow a non-stationary first-order Markov process. In summary, we consider a block model, where the feature values (and thus the corresponding random variable) remain constant within each block, and it changes from one block to the next following a first-order Markov process. The input load values are conditionally i.i.d. within each block, where the conditioning is on the feature values. To further reduce the complexity of the model (and ensure accuracy with limited data), we treat each day as being independent, and a group of days may be identical if their features are equal and can therefore be described by the same Markov process. Let k be the number of intervals within a day, each of them having i.i.d. realisations of the same random variable pair, and let $(\tilde{X}_i, \tilde{Y}_i)$ be the random variable pair corresponding to interval i. Fig. 1 illustrates our model for a single day. The transition probability matrix is estimated for $(\tilde{X}_1, \tilde{Y}_1)$ transitioning to $(\tilde{X}_2, \tilde{Y}_2)$ using the empirical distribution, *i.e.*, the histogram method (quantising the data and counting), which is a non-parametric method for probability estimation, as real measurement data do not necessarily follow specific distribution functions [13]. Let the distribution of $(\tilde{X}_1, \tilde{Y}_1)$ span 6 consecutive time slots, as in Fig. 1, and let (x_a, y_a) be the first realisation of $(\tilde{X}_2, \tilde{Y}_2)$. Each of the 6 realisations of $(\tilde{X}_1, \tilde{Y}_1)$ is treated as being a sample that could equally lead to state (x_a, y_a), resulting in 6 samples for estimating the probability $p(x_a, y_a | \tilde{X}_1, \tilde{Y}_1)$. While this may lead to an over-fit of the data, it is a compromise in order to handle the issue of having too few samples for estimating the transition probabilities. Moreover, this allows us to capture in our first-order Markov model some of the time correlation that spans multiple time slots, typically present in consumer loads, thereby implicitly modelling higher-order Markov processes. The right hand side of (4), which does not consider the stationarity assumption and which we denote I_{ub}, is approximated for a group of identical days (e.g., weekdays in summer), by:

$$\hat{I}_{ub} := \frac{1}{k}\left(\sum_{i=2}^{k} I(\tilde{X}_i, \tilde{X}_{i-1}; \tilde{Y}_i, \tilde{Y}_{i-1}) - \sum_{i=3}^{k} I(\tilde{X}_{i-1}; \tilde{Y}_{i-1})\right). \tag{6}$$

Note that the temporal average MI in (4) is replaced in (6) by the daily average MI as the days are considered to be identical within the same group. Let $I_{tr} := I(\tilde{X}_i, \tilde{X}_{i-1}; \tilde{Y}_i, \tilde{Y}_{i-1})$, we estimate the following:

$$I_{tr} = \sum_{\tilde{X}_i \in \mathcal{X}} \sum_{\tilde{X}_{i-1} \in \mathcal{X}} \sum_{\tilde{Y}_i \in \mathcal{Y}} \sum_{\tilde{Y}_{i-1} \in \mathcal{Y}} p(\tilde{x}_i, \tilde{x}_{i-1}, \tilde{y}_i, \tilde{y}_{i-1}) \times \log \frac{p(\tilde{x}_i, \tilde{x}_{i-1}, \tilde{y}_i, \tilde{y}_{i-1})}{p(\tilde{x}_i, \tilde{x}_{i-1}) p(\tilde{y}_i, \tilde{y}_{i-1})},$$

where $p(a)$ is estimated analogously to (3). It can be shown that $I_{ub} = \hat{I}_{ub}$ for a group of similar days if each time interval has exactly one realisation, though this may lead to a problem of data scarcity. This model should enable us to compute a better approximate for the actual value of the upper bound of the MI between the consumer and grid loads. Note also that the equality between the MI of the discrete (quantised) and continuous versions of $\tilde{X}_i$ and $\tilde{Y}_i$ is usually not achieved even if the MI is Riemann integrable [14], as data availability limits the number of quantisation levels.

It has been shown that for real SM data, e.g., the Irish Smart Meter Dataset [15], consumer demand, and implicitly its statistical distributions, changes according to a finite set of features. The authors of [16] show through the use of clustering techniques that the consumer demand in [15] changes as a function of the season, the day of the week, and the time of day. Hence, exploiting these features to model consumer load using a smaller finite set of random variables as proposed here represents a better approximation compared to assuming stationarity.

V. Numerical Experiments

Here, the proposed empirical MI approximation in (6) is compared with those in (2) and (5) through simulations using the Irish Smart Meter Dataset [15], as well as some synthetic load curves. In the rest of this section, we denote the values obtained using (6) by MI-v (time-*varying* Markov process), those from (5) as MI-s (*stationary* Markov process), and the values obtained using (2) as MI-i (*independent* and identically distributed).

A. Irish Smart Meter Dataset

In order to validate the proposed MI approximation method on real SM data, we simulate the MDPC and load-levelling controllers presented in [9] using data from Meter 1002 of the Irish Smart Meter Dataset [15] over a period of 450 days with an hourly time resolution. Different values of energy cost to privacy-loss are obtained by varying the relative price of privacy-loss, given in Rp/bit (100 Rappen (Rp) = 1 CHF), for both schemes. A brief description of both controllers are presented below, with the general simulation parameters given in Table I. The quantisation is equal for both X_τ and Y_τ, as shown in Table I, and is kept constant for the rest of this Subsection.

TABLE I
DEFAULT SIMULATION PARAMETERS.

Prediction Horizon, T:	12
MDPC Counting Window, N:	132
Number of $\mathcal{X}$ Bins, m:	20
Number of $\mathcal{Y}$ Bins, n:	20
Additive Smoothing, ε:	0.12
Reg. Coefficient, σ:	0.11
Battery Capacity:	6.4 kWh
Battery Power:	3.3 kW
Battery Efficiency, α:	96 %
Energy Price (high):	24.6 Rp/kWh
Energy Price (low):	13.15 Rp/kWh

1) MDPC: The MDPC scheme consists of solving at each time t the following optimisation problem

$$\begin{aligned} \underset{y,z}{\text{minimise}} \quad & \frac{1}{T+1}\sum_{\tau=t}^{t+T} c_\tau y_\tau + \mu\Phi(z) \\ \text{subject to} \quad & (y,z) \in \mathcal{F}_t, \end{aligned}$$

where T is the prediction horizon, c_τ is the price of energy at time τ, y_τ is the grid load, μ is the relative price of privacy loss, z is a vector of binary variables used in predicting the statistics, $\Phi(z)$ is an approximation of (2), and $\mathcal{F}_t$ enforces the system and binary constraints (see [9] for details).

2) Load-Levelling: The load-levelling scheme solves at each time t the optimisation problem

$$\begin{aligned} \underset{y}{\text{minimise}} \quad & \frac{1}{T+1}\sum_{\tau=t}^{t+T} c_\tau y_\tau + \frac{\mu}{T+1}\sum_{\tau=t}^{t+T}(y_\tau - y_{\tau-1})^2 \\ \text{subject to} \quad & y \in \bar{\mathcal{F}}_t, \end{aligned}$$

where $\bar{\mathcal{F}}_t$ enforces the system constraints (see [9] for details). This scheme attempts to flatten the grid load profile by penalising deviations from previous realisations, reducing private information leakage in the grid load.

According to the analysis performed in [16], the typical consumer load profile in the Irish Smart Meter Dataset [15] (excluding special days such as Christmas) can be described using the seasons (determined by standard equinox and solstice dates for Ireland), day of the week (weekdays and weekends), and time of day. Each day can be split into four consecutive time intervals as follows:

(a) *Overnight*: 10.30 p.m. to 6.30 a.m.
(b) *Morning*: 6.30 a.m. to 9.00 a.m.
(c) *Daytime*: 9.30 a.m. to 3.30 p.m.
(d) *Evening*: 3.30 p.m. to 10.30 p.m.

When computing MI-v, the realisations are grouped according to the features defined above, omitting the special days identified in [16]: 24th, 25th and 31st of December 2009, 1st, 9th and 10th of January 2010, and 4th of April 2010. Figures 2 and 3 illustrate the trade-off between the cost of energy (each point on the line resulting from a different μ value in the objective functions) and the three different MI approximations in autumn and spring, respectively; while Fig. 4 illustrates the trade-off for the whole simulation period, ignoring the seasonal changes in consumer demand. The MI-v curves shown in these figures are those for the weekdays. In autumn, comparable privacy protection can be achieved by both the MDPC and load levelling schemes when measured using MI-i, albeit at different energy costs. However, when one observes the resultant grid load curves for both schemes shown in Figures 5a and 5b, it is apparent that a clearer diurnal pattern is observable for the load levelling scheme. Hence, the MI-i measure is unable to capture the time correlation in the diurnal pattern. This time correlation is, however, reflected in the values of MI-v and MI-s, whose values indicate that the MDPC scheme is more private compared to the load levelling scheme, despite both schemes having similar MI-i values. This inability of MI-i to capture the time correlation, compared to MI-v and MI-s, is more apparent in spring as shown in Fig. 6a, where despite having a lower MI-i value of 0.147 bits, the load levelling scheme exhibits a clear diurnal pattern, while the MDPC scheme appears to be random and clearly more private despite having an MI-i value of 0.164 bits.

The difference between the lowest amount of privacy loss achieved in autumn and spring can be attributed to the higher load peaks in spring (6.90 kW max), which cannot be compensated for by the battery used in the setup. The battery is of sufficient power rating to compensate the load peaks in autumn, which has a maximum of 3.04 kW, enabling better privacy protection.

While the MI-v and MI-s approximations are able to capture some of the time correlation in the load levelling scheme, the MI-s method perceives the load-levelling scheme to be more private at higher energy costs as seen in Fig. 2b, eventually leading to values below that of MI-i, despite no significant changes in the grid load curves as the energy cost (and the weighting on privacy protection) increases (see Figures 5a and 5b for points LA1 and LA2), *i.e.*, no improvement in privacy protection. While the quantisation error does contribute to this reduction in MI-s, it is not a major factor as it is also present in MI-v, which does not show such a drastic improvement in privacy protection as the energy cost increases. By assuming stationarity, MI-s incorrectly models the feature-dependent distributions underlying both the LA1 and LA2 curves, leading to the observed reduction in privacy-loss despite both grid load curves being almost the same. Moreover, MI-s is also unable to capture the time correlation between realisations that are farther apart in time, e.g., the clear diurnal pattern in the LS2 curve, leading to a much lower MI estimate than expected for its clearly discernible periodic consumer consumption pattern seen in Fig. 6b. While the LS2 curve is flatter than the LS1 curve, it would be wrong to conclude that LS2 is more private than LS1, as privacy-leaking peaks have been masked in both, and both still exhibit a distinct diurnal pattern.

Both MI-v and MI-s estimates of the empirical MI are more accurate than MI-i. There is no significant difference when using MI-v and MI-s to assess the MDPC scheme, but the MI-s method is unable to accurately assess the load levelling scheme. Hence, among the three methods studied in this paper, MI-v has been shown to be the most accurate measure of

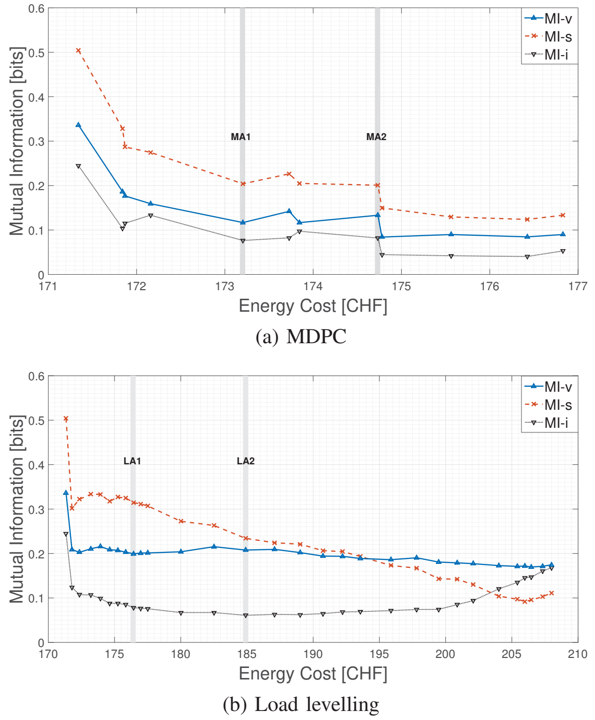

(a) MDPC

(b) Load levelling

Figure 2. Energy cost vs MI in autumn.

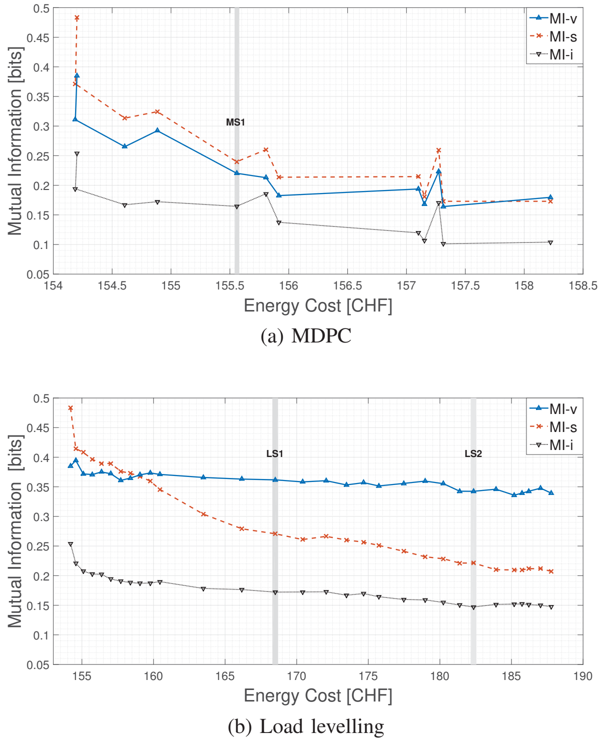

(a) MDPC

(b) Load levelling

Figure 3. Energy cost vs MI in spring.

empirical MI.

B. Synthetic Loads

The different methods of estimating empirical MI are further evaluated using a synthetic consumer load $\hat{X}$, and synthetic grid loads $\bar{Y}$ and $\hat{Y}$. The synthetic loads are quantised using 20

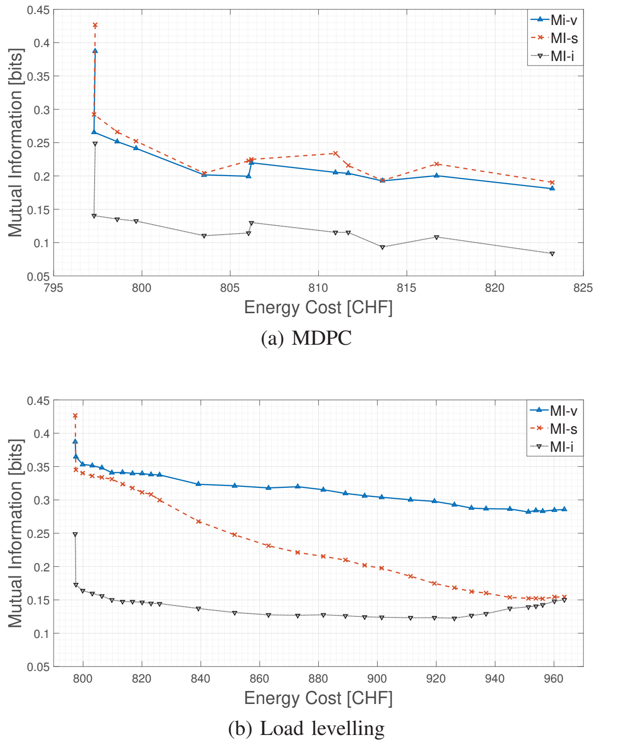

(a) MDPC

(b) Load levelling

Figure 4. Energy cost vs MI ignoring seasonality.

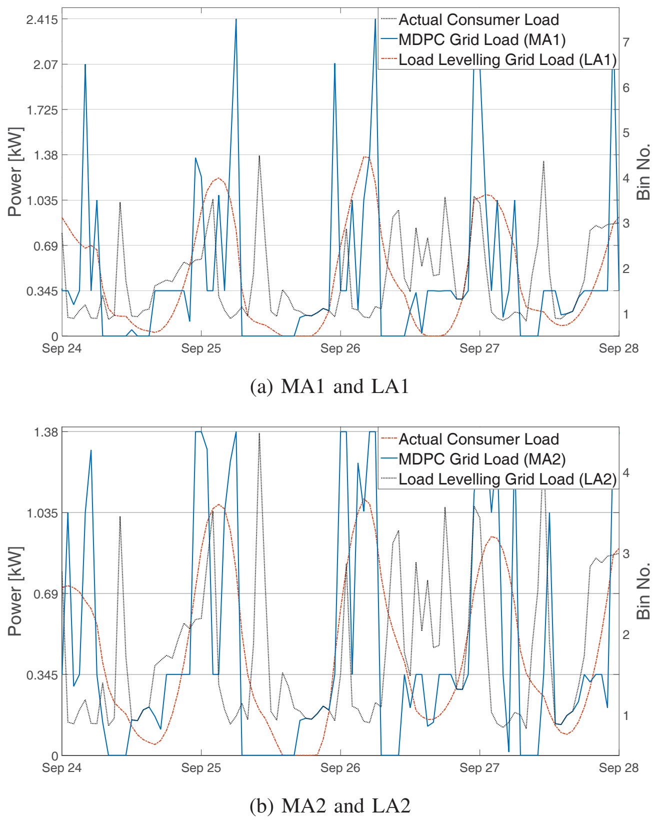

(a) MA1 and LA1

(b) MA2 and LA2

Figure 5. Consumer and grid load curves in autumn.

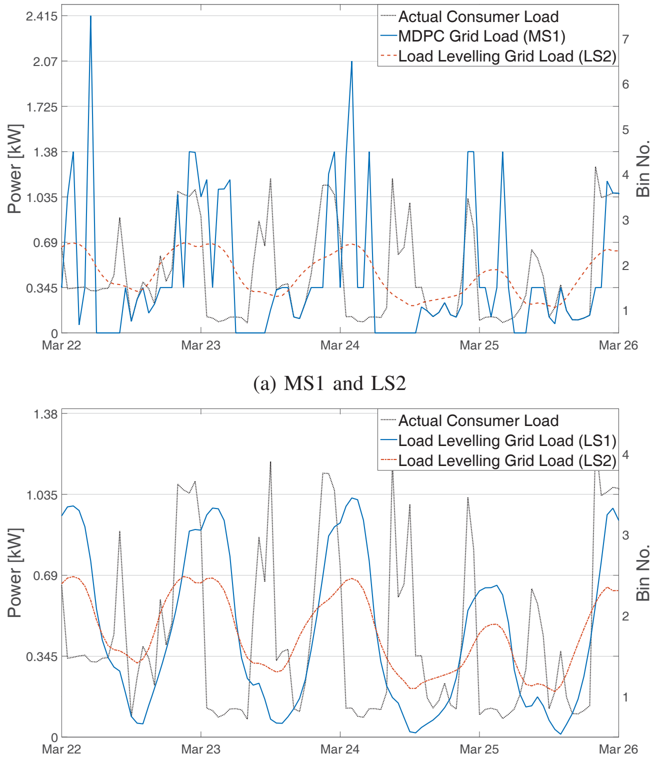

(a) MS1 and LS2

(b) LS1 and LS2

Figure 6. Consumer and grid load curves in spring.

bins, with even bin widths, and adjusted to match the range of each synthetic load curve. For the calculation of MI-v, every 24 consecutive realisations are treated as a single day, with four time intervals of equal length that define four different distributions. All days are treated as being identical.

1) Sinusoidal curves: The resultant grid load curves for the load levelling privacy protection scheme in Subsection V-A resemble smoothed out noisy sinusoids. To examine the accuracy of the proposed methods when used to assess sinusoidal curves, we calculate the empirical MI for two perfectly identical sinusoids $\hat{X} := \sin(B)+1$ and $\hat{Y} := \sin(B)+1$ (Case A), and for a sinusoid $\hat{X} := \sin(B) + 1$, and a normalised, phase-shifted sinusoid $\hat{Y}$ with Gaussian noise (Case B) given by

$$\hat{Y} := \frac{\sin(B+12)+\Lambda}{\max(\sin(B+12)+\Lambda)} + 1, \tag{7}$$

where $B = \{0, 6, 12, \dots\}$ is the angle in degrees, $\Lambda \sim \mathcal{N}(0, 0.05)$ is zero mean Gaussian noise, and $\max(A)$ computes the maximum value of A. The estimated empirical MI for these synthetic load curves are given in Table II.

None of the three methods accurately estimates the MI of two perfectly identical sinusoids, which in this example is theoretically 4.97 bits using the equation derived in [17], as there is no uncertainty when sinusoids are modelled using second-order Markov processes, and the number of $\mathcal{X}$ and $\mathcal{Y}$ bins are limited by the data available. When $\hat{Y}$ is a phase-shifted and noisy version of $\hat{X}$ as in Case B, MI-s shows an increase, contrary to the expected decrease due to the phase shift and added noise. This occurs as the $I(X_{t-1}; Y_{t-1})$ component in (5) decreases much faster than the $I(X_t, X_{t-1}; Y_t, Y_{t-1})$ component due to the stationary first-order Markov process assumption. MI-v does not exhibit this weakness due to its grouping of realisations into time intervals resulting in it implicitly modelling a higher-order Markov process. Additionally, as mentioned previously, modelling a sine wave as a second-order Markov process entails no randomness (given the realisations of X_{t-2}, and X_{t-1}, then $p_{X_t}(x) = 1$ for some x), and the MI in Case B would theoretically be due to the Gaussian noise alone. Hence, while all three empirical MI approximation methods fail to accurately capture the MI of sinusoids, MI-s performs the worst. This may help to explain the results seen in Fig. 2, where MI-s decreases with the increase in energy cost despite no noticeable change in the resultant grid-load.

TABLE II
EMPIRICAL MI FOR CASE A AND CASE B

	MI-v [bits]	MI-s [bits]	MI-i [bits]
Case A	2.83	0.985	3.99
Case B	2.03	1.40	2.09

2) Stepped grid load curves and the addition of Gaussian noise: Next, we analyse the MI approximation methods on stepped grid load curves $\bar{Y}$ and grid load curves $\hat{Y}_n$ that are increasingly noisier versions of the consumer load curve. The empirical MI for a single step (flat) grid load $\bar{Y}$, denoted $\bar{Y}_f$, and a two step grid load $\bar{Y}$, denoted $\bar{Y}_s$, are estimated. $\bar{Y}_f$ is the average of the consumer load $\hat{X}$, while $\bar{Y}_s$ step values are taken as the average consumer load $\hat{X}$ within the time interval where the step occurs. Mathematically, $\hat{Y}_n$ is given by

$$\hat{Y}_n := \frac{\hat{X}+\Omega}{\max(\hat{X}+\Omega)} + 1,$$

where $\Omega \sim \mathcal{N}(0, \sigma^2)$ is zero mean Gaussian noise, with variance σ^2 = 0.5α. Increasing α increases the Gaussian noise added to the consumer load.

Fig. 7a illustrates the original consumer load $\hat{X}$, $\bar{Y}_s$, and $\bar{Y}_f$, while Fig. 7b shows $\hat{Y}_n$ with $\alpha = 2$. Table III shows the empirical MI for the stepped grid load curves. A perfectly flat grid load curve gives zero empirical MI, and increasing the number of steps leads to more privacy-loss. The values of empirical MI for $\hat{X}$ and $\hat{Y}_n$ with increasing values of α are shown in Fig. 7c. All three methods of calculating empirical MI decrease with increasing noise. MI-i falls below the values of MI-v and MI-s as the $\hat{Y}_n$ curve increasingly loses its resemblance to $\hat{X}$, becoming increasingly Gaussian noise-like as seen in Fig. 7b for $\hat{Y}_n$ with $\alpha = 2$. At this noise level, MI-v and MI-s values are higher than MI-i as they capture the little remaining time correlation still present in the $\hat{Y}_n$ curve.

All three methods are capable of assessing the performance of some common privacy-protection methods, e.g., adding noise, heuristic load-levelling, and heuristic stepped control policies, with varying degrees of accuracy.

TABLE III
EMPIRICAL MI FOR $\bar{Y}_f$ AND $\bar{Y}_s$

	MI-v [bits]	MI-s [bits]	MI-i [bits]
$\bar{Y}_f$	0	0	0
$\bar{Y}_s$	0.721	0.244	0.967

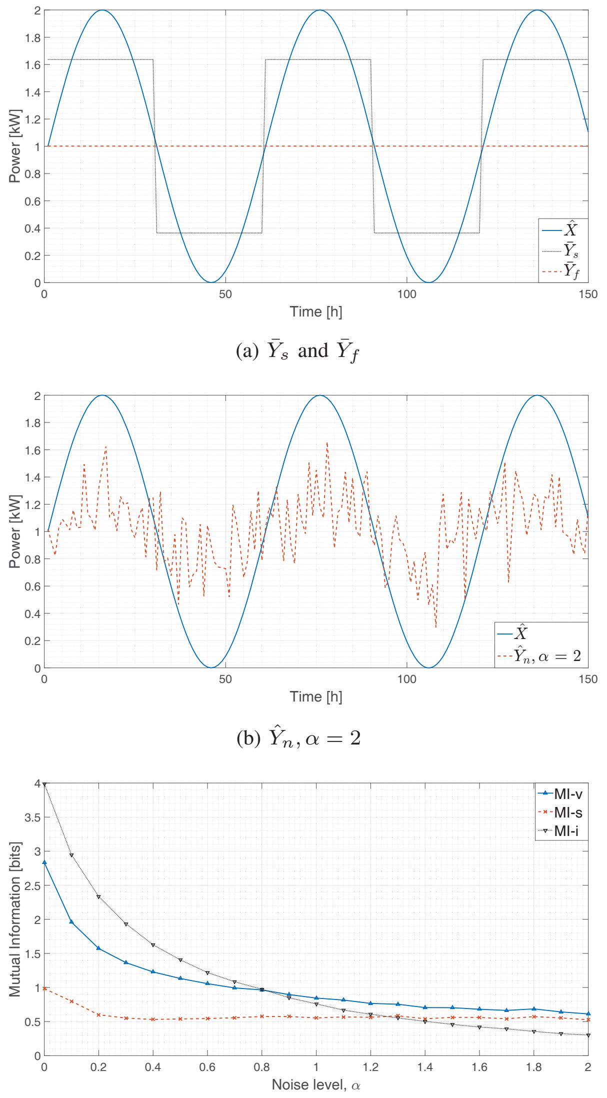

(a) $\bar{Y}_s$ and $\bar{Y}_f$

(b) $\hat{Y}_n, \alpha = 2$

(c) $\hat{Y}_n$ with increasing α

Figure 7. Results obtained by adding Gaussian noise and flattened load.

VI. Conclusions and Future Outlook

In this paper we calculated the empirical MI between consumer and grid loads by modelling their joint distribution and the consumer load distribution as feature-dependent first-order Markov processes. To reduce the number of random variables and avoid the issue of data scarcity, we grouped realisations according to features that define a change in consumer demand, e.g., time-of-day, day-of-week and season. This has also the advantage of implicitly modelling higher-order Markov processes. By means of numerical simulations, our formulation has been shown to produce a better approximation for the MI between consumer and grid loads. We remark that choosing an accurate method to approximate MI is crucial, as otherwise, privacy-loss may be underestimated.

Future research will focus on adapting the proposed MI approximation method in conjunction with a control policy in home energy management units, incorporating other features that define the random variables, using hidden Markov models for modelling the loads, and identifying other metrics to quantify the loss of consumer privacy.

References

[1] P. McDaniel and S. McLaughlin, "Security and privacy challenges in the smart grid," *IEEE Security and Privacy*, vol. 7, no. 3, pp. 75–77, 2009.

[2] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private memoirs of a smart meter," in *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, Zurich, Switzerland, 2010, pp. 61–66.

[3] PricewaterhouseCoopers, "Consumer intelligence series: Protect.me," PricewaterhouseCoopers, Tech. Rep., 2017. [Online]. Available: https://www.pwc.com/us/en/advisory-services/publications/consumer-intelligence-series/protect-me/cis-protect-me-findings.pdf

[4] G. Giaconi, D. Gündüz, and H. V. Poor, " Privacy-aware smart metering: Progress and challenges," unpublished. [Online]. Available: https://arxiv.org/pdf/1802.01166.pdf

[5] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proceedings of the 18th ACM conference on Computer and communications security (CCS '11)*, Chicago, Illinois, USA, 2011, pp. 87–98.

[6] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, "Minimizing private data disclosures in the smart grid," in *Proceedings of the 19th ACM conference on Computer and communications security (CCS '12)*, Raleigh, North Carolina, USA, 2012, p. 415.

[7] IRENA, "Electricity storage and renewables: costs and markets to 2030," International Renewable Energy Agency, Tech. Rep. October, 2017. [Online]. Available: http://www.irena.org/DocumentDownloads/Publications/IRENA_Electricity_Storage_Costs_2017.pdf

[8] O. Tan, J. Gomez-Vilardebo, and D. Gündüz, "Privacy-cost trade-offs in demand-side management with storage," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 6, pp. 1458–1469, 2017.

[9] J. Chin, T. Tinoco De Rubira, and G. Hug, "Privacy-protecting energy management unit through model-distribution predictive control," *IEEE Transactions on Smart Grid*, 2017, to be published.

[10] R. Malladi, D. H. Johnson, and B. Aazhang, "Data-driven estimation of mutual information between dependent data," unpublished. [Online]. Available: https://arxiv.org/pdf/1703.02468.pdf

[11] S. Gao, G. V. Steeg, and A. Galstyan, "Estimating mutual information by local Gaussian approximation," in *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence UAI'15*, Amsterdam, Netherlands, 2015, pp. 278–287.

[12] F. McLoughlin, A. Duffy, and M. Conlon, "The generation of domestic electricity load profiles through Markov chain modelling," *Euro-Asian Journal of Sustainable Energy Development Policy*, vol. 3, no. January - December, December 2010.

[13] D. Gross, P. Wiest, and K. Rudion, "Comparison of stochastic load profile modeling approaches for low voltage residential consumers," in *2017 IEEE Manchester PowerTech*, Manchester, UK, June 2017.

[14] T. M. Cover and J. A. Thomas, "Differential entropy," in *Elements of Information Theory*. Hoboken, NJ, USA: John Wiley & Sons, Inc., April 2005, ch. 8, pp. 243–259.

[15] Commission for Energy Regulation (CER), "CER smart metering project - electricity customer behaviour trial, 2009-2010," 2012. [Online]. Available: http://www.ucd.ie/issda/data/commissionforenergyregulationcer/

[16] S. Haben, C. Singleton, and P. Grindrod, "Analysis and clustering of residential customers energy behavioral demand using smart meter data," *IEEE Transactions on Smart Grid*, vol. 7, no. 1, pp. 136–144, 2016.

[17] J. M. Nichols, F. Bucholtz, and J. V. Michalowicz, "Calculation of entropy and mutual information for sinusoids," Naval Research Laboratory, Memo. Rep. NRL/MR/5650–09-9172, 2009. [Online]. Available: http://www.dtic.mil/dtic/tr/fulltext/u2/a494612.pdf