

RESEARCH

Open Access



The influence of differential privacy on short term electric load forecasting

Günther Eibl^{1*} , Kaibin Bao² , Philip-William Grassal³, Daniel Bernau³ and Hartmut Schmeck² 

From The 7th DACH+ Conference on Energy Informatics
Oldenburg, Germany. 11-12 October 2018

*Correspondence:

guenther.eibl@en-trust.at

¹Salzburg University of Applied Sciences, Center for Secure Energy Informatics, Urstein Süd 1, Puch/Salzburg, Austria

Full list of author information is available at the end of the article

Abstract

There has been a large number of contributions on privacy-preserving smart metering with Differential Privacy, addressing questions from actual enforcement at the smart meter to billing at the energy provider. However, exploitation is mostly limited to application of cryptographic security means between smart meters and energy providers. We illustrate along the use case of privacy preserving load forecasting that Differential Privacy is indeed a valuable addition that unlocks novel information flows for optimization. We show that (i) there are large differences in utility along three selected forecasting methods, (ii) energy providers can enjoy good utility especially under the linear regression benchmark model, and (iii) households can participate in privacy preserving load forecasting with an individual membership inference risk < 60%, only 10% over random guessing.

Keywords: Differential privacy, Load forecasting, Smart metering

Introduction

Smart metering data is said to be useful for improving the load forecasting task of energy providers (McDaniel and McLaughlin 2009; Li et al. 2010; Ilić et al. 2013; Bao and Lu 2015). With more accurate forecasts, energy providers gain an advantage for trading and scheduling electricity production and consumption ahead of time. Forecasting errors have to be compensated by buying control energy for stable electric grid operation. The highly volatile control energy prices charged for this compensation can be painful for the energy providers. In Germany of 2017, for example, the average control energy price was 49.67 EUR per MWh, but for 30 min, the price shot over 20,614.97 EUR per MWh.¹

On the other hand, monitoring electrical load from individual households incurs violation of privacy, as private behavior patterns are reflected in the energy consumption² (McDaniel and McLaughlin 2009; Molina-Markham et al. 2010; Lisovich et al. 2010). The amount of privacy violation varies depending on the monitoring time resolution of metering data (Eibl and Engel 2015). Using Differential Privacy (Dwork 2006) as privacy model, both time granularity and varying levels of the privacy parameters can be used to quantify and interpret the influence on privacy.

In addition to the privacy issue, the utility of individual data for load forecasting is naturally limited due to the stochasticity of domestic energy usage (Fan et al. 2009). To best of

our knowledge, no work exists that leverages individual (instead of aggregated) load data to gain a significant advantage on the domestic load forecasting task. This is why domestic load forecasting is performed using load data aggregated (added together) over large areas with many households.

In this paper, we investigate whether energy providers can benefit from aggregated smart metering data which is acquired in a privacy-friendly way. We formulate a privacy-preserving forecasting process that provides energy producers with forecast utility guarantees and households with strong, yet intuitive privacy guarantees, based on Differential Privacy. We make the following contributions:

- First time to regard energy provider's load forecasting task based on smart metering data with prescribed privacy guarantee,
- Practical design and evaluation of Differential Privacy for load forecasting, as well as comprehensible and interpretable calculation of membership inference risk using Differential Identifiability,
- Determination of the privacy-utility trade-off on real world data (Hong et al. 2014) using three realistic forecasting methods, and
- Demonstrating that differentially private load forecasting with a low membership inference risk $\rho < 0.6$ and strong utility is especially achievable under the linear regression benchmark model.

This paper is structured as follows. First, we introduce [preliminaries](#). We formulate our concept for realizing differentially private load forecasting in “[Differentially private metering and load forecasting](#)” section and present an evaluation in “[Experiments and results](#)” section. Finally, [related work](#) is presented before we [conclude with a discussion of practical implications](#).

Preliminaries

In the following, we provide [fundamentals regarding electricity grid metering](#), the underlying privacy model of this work in “[Differential Privacy](#)” section, and load forecasting approaches we use for electricity consumption prediction in “[Electric load forecasting methods](#)” section.

Electricity metering process (in Germany)

In this paper, we will discuss our problem setting in the context of the German metering and balancing process. Although the objective of metering for balancing in an electric power system is equal around the world, specific details in metering and settlement are subject to national and regional regulations. That is why we fix our process description to the well-documented German electrical power market. The relevant sources are the German electricity grid access regulation (StromNZV 2017), the German measuring point operation act (MsbG 2016) and the market rules for the implementation of balancing group accounting for electricity (MaBiS (Bundesnetzagentur 2011)).

In Europe, the electric grid is partitioned geographically into control areas, each operated by a transmission system operator (TSO). Every control area is subdivided into distribution grids operated by a distribution system operator (DSO). Transmission and distribution grid operators are government-regulated entities who are responsible for stable and reliable grid operation and non-discriminatory access to electricity production,

consumption and trading. To accomplish these two goals simultaneously, the TSO delegates the task of balancing supply and demand to the grid participants to some extent by charging the participants for any imbalance they cause. How imbalance is estimated and settled is subject to national regulations. (cf. (Commission Regulation 2017a; Commission Regulation 2017b; Federal Energy Regulatory Commission 2015))

Each control area is virtually partitioned into balancing groups which are basically time-dependent accounts for electric energy. An electricity customer (i.e., her grid connection point) is associated with exactly one balancing group which corresponds to the energy service provider and possibly to a specific tariff chosen by the customer. (cf. Sections 4, 5 StromNZV)

Before the roll-out of smart metering, residential electricity meters of customers with low or normal annual consumption were only read out annually or during the change of energy provider or tenant. Customers with an annual consumption above 100,000 kWh are subject to real-time load profile measurements which collect average and peak loads in quarter-hour intervals. With the roll-out of smart metering, additionally, customers with an annual consumption between 10,000 kWh and 100,000 kWh may be subject to load profile metering with quarter-hour resolution. (cf. Sections 55, 60 MsbG)

Figure 1 shows the essential roles and information flows as well as our envisioned privacy-preserving information flow in the metering and balance settlement process. Usually the TSO is the balancing group coordinator being responsible for determining the virtual balance of each balancing group and for charging them for imbalances. As the balancing group may be physically scattered among different distribution grids, the TSO needs to aggregate the information about the energy flows in the distribution grids from several DSOs. However, the DSO can not measure every grid connection point in real-time which is especially true for residential grid connection points. Therefore, the DSO estimates the residential loads either by using the synthetic or the analytical method (Step 1 and 2 in Fig. 1). The synthetic method uses parameterized standard load profiles which are scaled by a forecasted annual energy consumption of each customer. For the analytical method, the DSO subtracts the real-time metered load profiles and estimated transmission losses from the overall load profile of its distribution grid. The remainder is the load profile of the non-metered residential grid connection points, which is then attributed according to a forecasted annual energy consumption of each customer (Step 3 in Fig. 1). (cf. Section 1.2 MaBiS (Bundesnetzagentur 2011) and Section 3.8 of the Distribution Code 2007 (Epe et al. 2007))

The TSO finally aggregates the load profiles from all DSOs to determine the load profile of each balancing group (Step 4 in Fig. 1). This overall ex-post balance in each group is used to settle the costs for the actual imbalance during the grid operation. The party responsible for the balancing group whose imbalance helps compensating the overall grid imbalance, is being paid for the grid support. All balancing group parties receive their corresponding load and balance measurements in order to retrace the bill and to improve on the load predictions (Step 5 in Fig. 1) (cf. Section 2 MaBiS). We envision that after the smart meter rollout, the aggregate consumption for each zone of each balancing group can be obtained using a privacy mechanism (Step (2b)) and used for improving the forecasting of the balance group.

Technically, the current metering process is not differentially private as the aggregated load of a balancing group is not perturbed. However, the current metering process

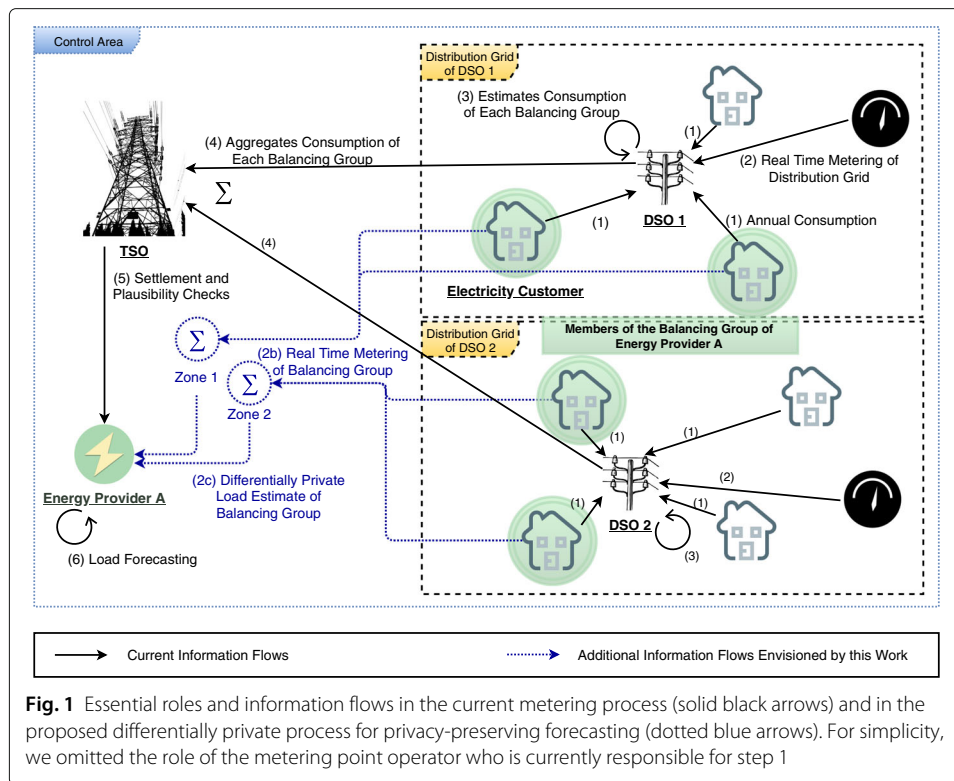


Fig. 1 Essential roles and information flows in the current metering process (solid black arrows) and in the proposed differentially private process for privacy-preserving forecasting (dotted blue arrows). For simplicity, we omitted the role of the metering point operator who is currently responsible for step 1

based on non-smart meters is generally not considered as serious privacy violation since residential electricity measurements are read out only once per year.

Differential Privacy

Differential Privacy, originally proposed by (Dwork 2006), is the current gold standard for data privacy. It is achieved by perturbing the result of a query function $f(\cdot)$ s.t. it is no longer possible to confidently predict whether the result was obtained by querying data set D_1 or some other data set D_2 differing in one individual. Thus, privacy is provided to each participant in the data set as their presence or absence becomes almost negligible for computing perturbed query results. To inject noise into the result of some arbitrary query $f(\cdot)$, mechanisms K_f are utilized. Mechanisms add noise sampled from a probability distribution to $f(\cdot)$ and are differentially private if they fulfill Definition 1.

Definition 1 (Differential Privacy). A mechanism

$K_f : DOM \rightarrow R$ is (ϵ, δ) -differentially private if for all data sets $D_1, D_2 \subset DOM$ differing in only one individual and for all possible outputs $S \subseteq R$:

$$Pr[K_f(D_1) \in S] \leq e^\epsilon * Pr[K_f(D_2) \in S] + \delta \quad (1)$$

The additive δ is interpreted as the probability of protection failure and required to be negligibly small $\approx \frac{1}{|D_1|}$. We refer to (Dwork and Roth 2014) for the proof. Another commonly used, more strict definition calls a mechanism ϵ -differentially private if it is $(\epsilon, 0)$ -differentially private. Differential Privacy has the appealing property that it holds independent of any side knowledge of the adversary. Therefore, an adversary may know everything but not whether S was computed using D_1 or D_2 . We call a data set differentially private if it has been obtained by a differentially private mechanism.

The query is further specified as a series of k identical aggregate queries f_i with co-domain $R = \mathbb{R}$ each. The added noise must hide the influence of any individual in the original result of the composed query $f = (f_1, \dots, f_k)$. The maximum influence of an individual on $f(\cdot)$ is the *global sensitivity* $\Delta f = \max_{D_1, D_2} \|f(D_1) - f(D_2)\|_1$.

A popular mechanism for perturbing the outcome of numerical query functions is the *Laplace mechanism*, proposed by (Dwork et al. 2006). It adds noise calibrated w. r. t. the global sensitivity by drawing a random sample from the Laplace distribution with mean $\mu = 0$, scale $\lambda = \frac{\Delta f}{\epsilon}$ according to Theorem 1.

Theorem 1 (Laplace Mechanism). *Given a series of k identical numerical query functions $f = (f_1, \dots, f_k) \in \mathbb{R}^k$, the Laplace Mechanism*

$$K_{Lap}(D, f, \epsilon) := f(D) + (z_1, \dots, z_k) \tag{2}$$

is an $(\epsilon, 0)$ -differentially private mechanism when all z_i with $1 \leq i \leq k$ are independently drawn from the random variable $Z \sim Lap\left(z, \frac{\Delta f}{\epsilon}, 0\right)$.

Again, for proof, we refer to (Dwork et al. 2006). To apply Theorem 1 to smart metering, i.e., a distributed setting, we use the gamma distribution suggested for distributed noise generation by Ács et al. (Ács and Castelluccia 2011). The following Lemma 1 leads to the generation of gamma noise that satisfies the Laplace mechanism. We use this divisibility to formulate a distributed differentially private metering process in “Differentially private metering process” section.

Lemma 1 (Divisibility of Laplace distribution (Kotz et al. 2001; Ács and Castelluccia 2011)). *Let $Z(\lambda)$ denote a random variable from a Laplace distribution with density $f(x, \lambda) = \frac{1}{2\lambda} e^{-\frac{|x|}{\lambda}}$. Then the distribution of $Z(\lambda)$ is infinitely divisible. This means that for every integer $n \geq 1$ it can be represented as a sum of n random variables $Z(\lambda) = \sum_{i=1}^n X_i$. Here, each $X_i = G_1(n, \lambda) - G_2(n, \lambda)$. $G_1(n, \lambda)$ and $G_2(n, \lambda)$ are i. i. d. random variables having gamma distribution with density $g(x, n, \lambda) = \frac{(1/\lambda)^{1/n}}{\Gamma(1/n)} x^{n-1} e^{-x/\lambda}$ defined for $x \geq 0$.*

When a function is evaluated multiple times an overall *privacy loss* occurs. Under worst case assumptions, the *sequential composition theorem* of Differential Privacy states that a series of k evaluations of any (ϵ, δ) -differentially private mechanism K_f on the same set of individuals results in $(k\epsilon, k\delta)$ -Differential Privacy. However, recent results by Dwork et al. (Dwork et al. 2010) and Kairouz et al. (Kairouz et al. 2017) prove that actually sub-linear increases in ϵ are achieved under k -fold composition when allowing a small $\tilde{\delta}$ under Theorem 2.

Theorem 2 (k -Fold Adaptive Composition for Homogeneous Mechanisms). *For any $\epsilon > 0$ and $\delta \in [0, 1]$, and $\tilde{\delta} \in (0, 1]$ the class of (ϵ, δ) -differentially private mechanisms satisfies $(\tilde{\epsilon}_{\tilde{\delta}}, 1 - (1 - \delta)^k (1 - \tilde{\delta}))$ -Differential Privacy under k -fold adaptive composition, for*

$$\tilde{\epsilon}_{\tilde{\delta}} = \min \left\{ \begin{array}{l} k\epsilon \\ \frac{(e^\epsilon - 1)k\epsilon}{e^\epsilon + 1} + \epsilon \sqrt{2k \ln \left(e + \frac{\sqrt{k\epsilon^2}}{\tilde{\delta}} \right)} \\ \frac{(e^\epsilon - 1)k\epsilon}{e^\epsilon + 1} + \epsilon \sqrt{2k \ln \left(\frac{1}{\tilde{\delta}} \right)} \end{array} \right. . \tag{3}$$

When operating in high privacy regimes ($\epsilon \ll 1$), the term $\frac{(e^\epsilon - 1)k\epsilon}{e^\epsilon + 1} \approx k\epsilon^2$ illustrates the sub-linear loss of privacy under k -fold composition. Even though composition allows to determine the privacy decay by growth in ϵ over a series of queries, a rational explanation for the actual choice of ϵ is missing. To the best of our knowledge, there is no approach

for giving concrete guidance on choosing ϵ . Nonetheless, we are convinced that providing a more comprehensible interpretation of ϵ and the corresponding guarantee is crucial for acceptance of Differential Privacy in practice.

Consequently, we apply a *belief model* in this work to give smart metering users a better understanding of their protection guarantee ϵ . The foundation of this model led (Lee and Clifton 2012) to define *Differential Identifiability*, a privacy notion slightly differing from *Differential Privacy*. For convenience, we restate the definition of Differential Identifiability in Definition 2.

Definition 2 (Differential Identifiability). *Given an original data set D , a randomized mechanism K satisfies ρ -Differential Identifiability if among all possible databases D_1, D_2, \dots, D_m differing in one individual w. r. t. D the posterior belief P , after getting the response r , is bounded by ρ :*

$$P(D_i | K(D) = r) \leq \rho \quad . \quad (4)$$

ρ -Differential Identifiability implies that after receiving a mechanism's output r the true data set D can be identified by an adversary with confidence $\leq \rho$. Findings by (Li et al. 2013) show that Differential Privacy and Differential Identifiability are actually equal when $m = 2$ since Differential Privacy considers only two neighboring data sets D_1, D_2 by definition. If this condition is met, according to (Li et al. 2013), the relation between ρ and ϵ is:

$$\epsilon = \ln \left(\frac{\rho}{1 - \rho} \right) \quad \text{and} \quad \rho = \frac{1}{1 + e^{-\epsilon}} > \frac{1}{2} \quad . \quad (5)$$

Consequently, the confidence ρ provides a simplified interpretation of the actual membership inference risk when applying $(\epsilon, 0)$ -Differential Privacy. When $\delta > 0$, we define that the confidence of ρ holds with probability $1 - \delta$. We use this method to substantiate our results in "[Application of differential privacy](#)" section.

Electric load forecasting methods

Three different forecasting methods are used within this work. The first method is the [benchmarking forecasting model](#) of the 2012 Global Energy Forecast Competition (GEFCom 2012). The other two methods, [CountingLab's forecasting method](#) and [Lloyd's forecasting method](#), were the two highest ranked forecasting methods of the competition. For the first time, the impact of Laplacian noise for differential privacy on realistic forecasting methods is studied ("[Experiments and results](#)" section).

Global energy forecast competition 2012 (GEFCom 2012)

In 2012, an electric load forecasting competition (GEFCom 2012) (Hong et al. 2014) was conducted. Here, the time span of the given historical load data of an ISO in the USA was approximately 4.5 years in hourly readout intervals from 20 zones. Additionally, historical temperature data of 11 nearby weather stations were given, but there was no information about the association between weather stations and zones. For the forecasting time period, the temperature data was not given and needed to be forecasted, too. A limited amount of tuning was possible due to allowing multiple submissions and directly showing the resulting score.

The statistics of the historical load data are plotted in Fig. 2. Zone 4 is the smallest zone with a mean load of only 0.575 MW. In the right panel it can be seen that Zone 9 exhibits outliers with low consumption values which indicates metering issues or local blackouts.

Benchmark forecasting model of GEFCom 2012

Hong (Hong et al. 2014) provided a linear regression model as a benchmark for GEFCom 2012 competition. A linear regression model for load forecasting has the general form

$$F_{z,t} = \beta_0 + \sum_{j=1}^p \beta_j x_{t,j} + e_t \quad (6)$$

where $F_{z,t}$ is the forecast of the aggregate energy consumption of zone z in time slot t , β_j are the parameters of the model, $x_{t,j}$ are the independent variables, e_t is the residual error which cannot be explained by the model.

The 20 benchmark models (one per zone) consider a total of $p = 313$ explanatory temperature and calendar variables $x_{t,j}$ or cross-effects which are described in the GEFCom 2012 paper (Hong et al. 2014). For each zone, only one of the 11 weather stations is chosen to provide the temperature values for the linear model. The choice was made by fitting the model for each of the weather stations and choosing the one with the smallest training error $\sum_{t=1}^k |e(t)|$ where the errors are summed up for all time points. Using that final model, a forecast for the week following the given historical data was to be estimated. While the explanatory calendar variables can be easily obtained, no forecast for temperature T_s of the weather stations was given. The benchmark model constructed temperature forecasts by “averaging the temperature at the same date and hour over the past four years” (Hong et al. 2014).

CountingLab’s forecasting method

Among the forecasting methods referenced in the GEFCom 2012 paper (Hong et al. 2014), CountingLab (Charlton and Singleton 2014) achieved the best test score in the competition. As the benchmark model, it relies on multiple linear regression (6). However, in contrast to the benchmark model, not 20 single forecasts are obtained (one per zone) but 3,840 forecasts $F_{z,h,S,w}$ for each combination of zone z , hour of the day h , season S and day type. As a benefit the number of independent variables is much smaller than for the benchmark model: only nine parameters (interactions of temperature, day number and day number within the season) must be fit per linear model.

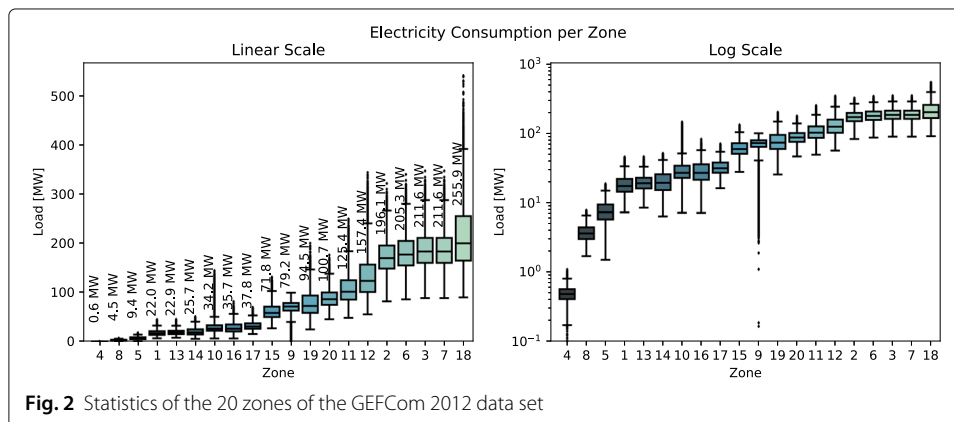


Fig. 2 Statistics of the 20 zones of the GEFCom 2012 data set

The needed temperature forecast is the mean of historical temperatures. In order to win the competition the authors spent additional effort, see (Charlton and Singleton 2014) for more details.

Lloyd's forecasting method

Lloyd (Lloyd 2014) achieved the second best test score in the competition. First, the temperature was estimated as the sum of a smooth trend and a daily periodic estimate using Gaussian processes with squared exponential and periodic kernels, respectively.

The prediction is a weighted ensemble of three forecasting methods: (i) the benchmark model (see “[Benchmark forecasting model of GEFCom 2012](#)” section) with weight 0.1, (ii) a gradient boosting machine (Friedman 2001) with weight 0.765 and (iii) a Gaussian process regression with weight 0.135. The weights have been chosen by manual tuning.

For each zone, a separate boosting model was learned using as input the time of day $\in [0, 1]$, the time within the week $\in [0, 7]$, the temperature predictions and smoothed temperature predictions of all weather stations. Note that the loads are not used as inputs, only as response values.

The third method uses Gaussian process regression with three additive kernels for forecasting that all depend on time: two squared exponential kernels should explain the variation of the load by two different length scales; the third, periodic kernel should model the periodic behavior.

Differentially private metering and load forecasting

As we have discussed in “[Electricity metering process \(in Germany\)](#)” section, the energy provider as a responsible party for a balancing group has the incentive to forecast the aggregate load of her customers with low error so that she can trade or schedule energy production ahead of time for lower costs. Smart metering could provide the means to realize a more accurate forecast by using load monitoring of individual households. However, monitoring individual loads conflicts with customer privacy interests. Due to the high stochasticity of individual loads, proficiently grouping households based on geographic or topological areas (zones) is beneficial to the forecasting performance (cf. (Fan et al. 2009)). Thus, monitoring *individual* loads may not have benefits anyway.

Therefore, we propose that the energy provider and the customers split the difference by agreeing on a trade-off between forecasting accuracy and customer privacy. For that, Differential Privacy is guaranteed for the customer by grouping households into zones and applying the Laplace mechanism on the *aggregated* load of each zone. Individual loads do not need to be disclosed, since in this paper we assume the usage of a privacy-preserving protocol for the smart metering infrastructure based on additional homomorphic encryption (Li et al. 2010; Erkin and Tsudik 2012) or masking (Acs and Castelluccia 2011; Knirsch et al. 2018). These protocols enable the calculation of the sum of all the household's load values of a zone at each time point without providing the individual values.

Since the load of the whole balancing group used for balance compensation must be provided by the TSO, this could be used for a direct forecast of the aggregate of the balancing group. Smart metering data are only useful, if better forecasts can be obtained. Possibly, forecasting could be improved by providing differentially private, *zonal* sub-aggregates of the region obtained by the DSOs. As shown in “[Experiments and results](#)”

section depending on the forecasting method, using this so-called hierarchical forecast compared to the direct forecast may not even improve the forecasting performance.

If there is an advantage using the hierarchical forecast, the acquisition of smart metering data has to satisfy the privacy interests of the affected customers. For each forecasting method, the limit for the privacy level of the differentially private method depends on the additional error introduced by it. If the forecasting error exceeds the direct forecasting error, smart metering data do not help the energy provider for trading or scheduling tasks. In “[Basic Forecasting Problem](#)” section and “[Evaluation of Forecasting](#)” section, we reflect that idea in the definition of the load forecasting problem and the utility definition.

We envision that the energy provider offers different energy tariffs coupled with a specific privacy protection level in terms of different λ values for the Laplace mechanism applied to the calculation of the zonal aggregate. The differentially private metering process is detailed in “[Differentially private metering process](#)” section. The customer interprets the privacy level coupled with the tariff by deriving her own effective privacy protection in terms of Differential Identifiability. Using that assessment, she can perform an informed decision about her energy tariff and how much privacy she is willing to trade. This is described in “[Application of differential privacy](#)” section.

Basic Forecasting Problem

We consider a control area, called region for simplicity, that is divided into Z zones containing n_z households. All households i of a zone z provide their load measurements $l_{z,i,t'}$ at several time points t' . The zone aggregators (DSOs) calculate the sum of all the household's load values $L_{z,t'}$ at each time point t' without receiving the individual values. Therefore, for each zone for each time point t' , only the sum $L_{z,t'}$ of the load values $l_{z,i,t'}$ of all the households i of the zone is available. These zonal aggregate loads are available at past time points t'_1, \dots, t'_k . The goal then is to predict the regional aggregate load $\mathcal{L}_{t'}$ which is the sum of the zone's loads:

$$L_{z,t'} := \sum_{i=1}^{n_z} l_{z,i,t'} \quad , \quad \mathcal{L}_{t'} := \sum_{z=1}^Z L_{z,t'} \quad . \quad (7)$$

Based on values available at times t'_1, \dots, t'_k the forecasting problem consists of producing forecasts F_{t_1}, \dots, F_{t_T} for the regional aggregate load \mathcal{L} for a sequence of forecast horizons t_1, \dots, t_T in the future ($t_1 > t'_k$). For forecasting, not only past aggregate load values are available but also additional factors $\vec{x}_{t'_1}, \dots, \vec{x}_{t'_k}$, where vector $\vec{x}_{t'} = (x_{t',1}, \dots, x_{t',p})$ summarizes p different explanatory variables. Typical information that is summarized in $\vec{x}_{t'}$ is, for example, the temperature, the hour of the day or the season (cf. “[Electric load forecasting methods](#)” section).

Direct forecasting

Two variants are distinguished. The first variant is called *direct* forecasting that predicts each $\mathcal{L}_{t'}$ of the prediction period based on the past regional aggregate values $\mathcal{L}_{t'_1}, \dots, \mathcal{L}_{t'_k}$ and other factors $\vec{x}_{t'_1}, \dots, \vec{x}_{t'_k}$. Note that no zonal aggregate loads are available for forecasting. The corresponding forecast value is denoted by $F_{\text{direct},t}$.

Hierarchical forecasting

The second variant is called *hierarchical* forecasting that is allowed to use the zonal aggregate loads. First, the aggregate load per zone is predicted and the prediction of the region is obtained by the sum of the predicted zonal aggregates. For each zone z , based on the past values $L_{z,t'_1}, \dots, L_{z,t'_k}$ and other factors $\vec{x}_{t'_1}, \dots, \vec{x}_{t'_k}$, future values $L_{z,t_1}, \dots, L_{z,t_T}$ are forecasted. With the forecasts of the zonal aggregates denoted as $F_{z,t_1}, \dots, F_{z,t_T}$, the overall aggregate \mathcal{L} is then estimated

$$\mathcal{F}_t = \sum_{z=1}^Z F_{z,t} \quad . \quad (8)$$

Hierarchical forecasting is assumed to be beneficial when the loads of different zones must be predicted differently depending on the forecast inputs. For example, loads exhibit more or less distinctive maximum values also at different times of the day. The temperature might as well be described more accurately for smaller, more homogeneous zones.

Hierarchical forecasting with differentially private, perturbed data is effectively the same problem as before with the difference that for each zone instead of the exact aggregates $L_{z,t'_1}, \dots, L_{z,t'_k}$ solely perturbed aggregates $\hat{L}_{z,t'_1}, \dots, \hat{L}_{z,t'_k}$ are used. Since resulting forecast of the regional aggregate then depends on the amount of added noise, it is denoted as F^λ .

Evaluation of Forecasting

As stated before, we need to add noise to the original data to achieve Differential Privacy. The noise added to the aggregate of each zone yields to the Differential Privacy property of the aggregate load data of each zone z . Due to the immunity against post-processing (cf. “Differential Privacy” section), the overall aggregate load that is computed as the sum is differentially private.

It is clear that in a practical setting, perturbed data can only be of use if noise does not destroy the prediction performance of the data. For calculating utility, we compare the direct, non-hierarchical forecast with the hierarchical forecast using differentially private load data of the zones.

The utility of load forecasts of period t_1, \dots, t_T is assessed by two error measures. Firstly, the commonly used mean absolute percent error *MAPE*, a scale-free error measure that enables a comparison of our results with results for other datasets. Secondly, the mean absolute error *MAE* that allows to compare different forecasting methods for the same (GEFCom) data. Both error measures are computed according to their names, i.e.,

$$MAE := \frac{1}{T} \sum_{t=t_1}^{t_T} |F_t - L_t| \quad , \quad (9)$$

$$MAPE := \frac{1}{T} \sum_{t=t_1}^{t_T} \left| \frac{F_t - L_t}{L_t} \right| \quad . \quad (10)$$

This way the error can be assessed both for forecasting the aggregate load L_z of a zone z and the overall aggregate load \mathcal{L} .

We define utility u^λ as the relative gain we achieve by switching from non-hierarchical to hierarchical forecast with perturbed data,

$$u^\lambda := \frac{MAE_{\text{direct}} - MAE^\lambda}{MAE_{\text{direct}}} \quad (11)$$

The error measure MAE^λ uses the forecasts F^λ of the hierarchical forecast with perturbed data for the overall aggregate load \mathcal{L} , where MAE_{direct} is the error of the direct load forecast for the overall aggregate load \mathcal{L} . Since the regional aggregate values are known in any case, direct forecasting can always be applied. Therefore, hierarchical forecasting only makes sense if it performs better than direct forecasting. Consequently, when the perturbation factor λ gets too large, it causes the error MAE^λ to exceed the direct forecasting error MAE_{direct} and the utility becomes negative $u^\lambda \leq 0$.

Differentially private metering process

In this work, we strive to bring energy providers in the position to train load forecast models on differentially private aggregated data from electricity customers. Differential Privacy is required due to possible insufficiencies of pure aggregation for privacy protection (Dwork 2006; Buescher et al. 2017).

As stated previously, energy providers desire to limit deviation of their forecasting algorithms due to differentially private noise added to load forecasting training data by specifying an upper bound for acceptable forecasting errors. In turn, this will lead to an upper bound for acceptable noise scales λ which is needed according to the Laplace mechanism for achieving Differential Privacy. More specifically, by Theorem 1 each household is provided $\frac{\Delta f}{\lambda} = \epsilon$ -Differential Privacy. The application of the Laplace mechanism results in three challenges in the scenario of this paper.

Firstly, we do not want to have perturbation $L_{z,t'_1}, \dots, L_{z,t'_k} \rightarrow \hat{L}_{z,t'_1}, \dots, \hat{L}_{z,t'_k}$ done by the energy provider to avoid assumptions about trustworthiness. Instead, we desire perturbation to be performed at the data sources directly, i. e. a smart meter adds noise itself for each point in time t' . Following Lemma 1, we realize this by decomposing Laplace noise into gamma noise for distributed noise generation at household level as stated in Eq. 12. The provider has to compute the sum for each zone to obtain the noisy total consumption, see Eq. 13.

$$\hat{l}_{z,i,t'} = l_{z,i,t'} + (\mathcal{G}_1(n_z, \lambda) - \mathcal{G}_2(n_z, \lambda)) \quad (12)$$

$$\hat{L}_{z,t'} = \sum_{i=1}^{n_z} \hat{l}_{z,i,t'} = L_{z,t'} + \mathcal{Z}(\lambda) \quad (13)$$

Secondly, the training data is represented by a time series t'_1, \dots, t'_k of each electricity customer's energy data, i.e., involving always the same set of households. Consequently, privacy decays over time as more information is revealed. For measuring the accumulated privacy loss, we apply Theorem 2 to obtain the total privacy loss $\tilde{\epsilon}_\delta$ as a function of ϵ, δ and time k .

Thirdly, the accumulated privacy guarantee, $\tilde{\epsilon}_\delta$, is hard to interpret for consumers (i.e., electricity customers). Our envisioned process addresses this by translating $\tilde{\epsilon}_\delta$ into an interpretable risk ρ by Eq. 5. ρ represents the upper bound for the confidence of an adversary trying to infer the membership of a single household in $\hat{L}_{t',z}$. We have almost perfect privacy if an attacker is unable to confidently distinguish whether a household contributed

to the sum or not, i.e., $\rho \approx 0.5$ (random guessing). In contrast, if $\rho \approx 1$ the privacy level is extremely low. To provide a reasonably good protection, we aim to bound the confidence at $\rho = 0.6$, meaning that even in worst case situations an adversary is not able to identify that a household contributed with more than 60% confidence.

Our process of applying Differential Privacy has several benefits. The energy provider does not have to perform any perturbation as noise is added locally by each meter and adds up to noise following the Laplace mechanism. In addition, providers can select the amount of noise λ they tolerate with regard to their forecasting algorithms. λ then gets propagated to households who resolve it to their corresponding ρ to see how much data privacy the energy provider actually ensures.

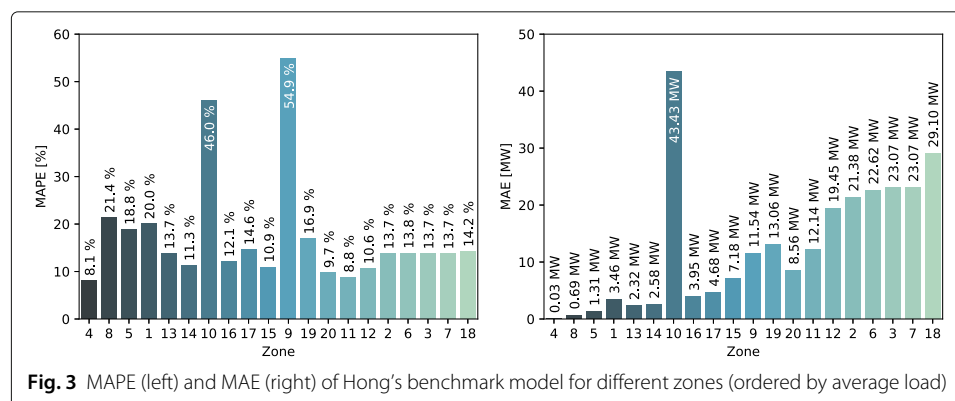
Experiments and results

In this section, three different models for forecasting the [GECOM data set](#) are trained. After confirming the correctness of the implementations by applying the forecasts to unperturbed data, the sensitivity of the forecasting performance on Laplacian noise of different scales λ is assessed. As the noise scale λ does not lend itself to describe the achieved privacy in a comprehensive way, such a description is developed in “[Application of differential privacy](#)” section based on the Differential Identifiability notion. Using all of the above, the privacy-utility trade-off will be described.

Forecast results

We re-implemented Hong’s linear regression benchmark model³ and CountingLab’s forecast model. For sake of simplicity, we omitted 2 of the improvements of CountingLab’s model. Lloyd’s method did not need to be implemented because the source code is freely available⁴. Only adaptations facilitating the handling of many different input files have been necessary.

Firstly, we verified the correctness of the implementation for unperturbed data. The MAPE and MAE of the non-perturbed forecast by Hong’s model for each zone are depicted in Fig. 3. The zones are sorted by their average load from left to right. Zone 9 and 10 have prominently high errors. As Fig. 2 shows, the outliers in Zone 9 indicate metering errors or power outages. In Zone 10, the average monthly consumption suddenly tripled starting in January 2008, indicating a change of the grid configuration. As the forecasting time period is after January 2008, this may be the cause of the high forecasting errors.



Similarly, both CountingLab's and Lloyd's forecast models are inaccurate for zones 9 and 10. However, for both of these models, the averaged errors for the remaining zones are smaller compared to the benchmark model. A comparison between unperturbed direct forecast and the unperturbed hierarchical forecast shows that the hierarchical forecast for the benchmark method lowers the average error by 12 MW. This results in a utility of 7.8% (first line of Table 1) and means that our privacy mechanism should not introduce additional errors much above 12 MW in order to avoid too negative utilities. Surprisingly, the hierarchical forecast is worse than the direct forecast for the other two models resulting in negative utilities.

Now, the impact of varying levels of noise on the forecast performance is evaluated. Figure 4a shows the forecasting error of perturbed hierarchical forecasting of Hong's benchmark model using increasing levels of perturbation. We train the models and run the forecast 10 times each with different random seeds. In some cases, the error even decreases. The red line indicates our utility-limit of 12 MW above unperturbed error (blue line). With $\lambda = 56, 234$, all runs still stay below this limit. Starting at $\lambda = 100,000$, some runs start to show higher error than the unperturbed direct forecast.

While the performance of CountingLab's forecast is better than the benchmark model for unperturbed data, the performance is highly negatively affected by the noise. This can be seen in Fig. 4b where the MAE quickly rises with λ .

The main difference between CountingLab's method and the benchmark model lies in the construction of many small models that use a smaller amount of data, each. It seems plausible that noise has a greater negative effect on such approaches.

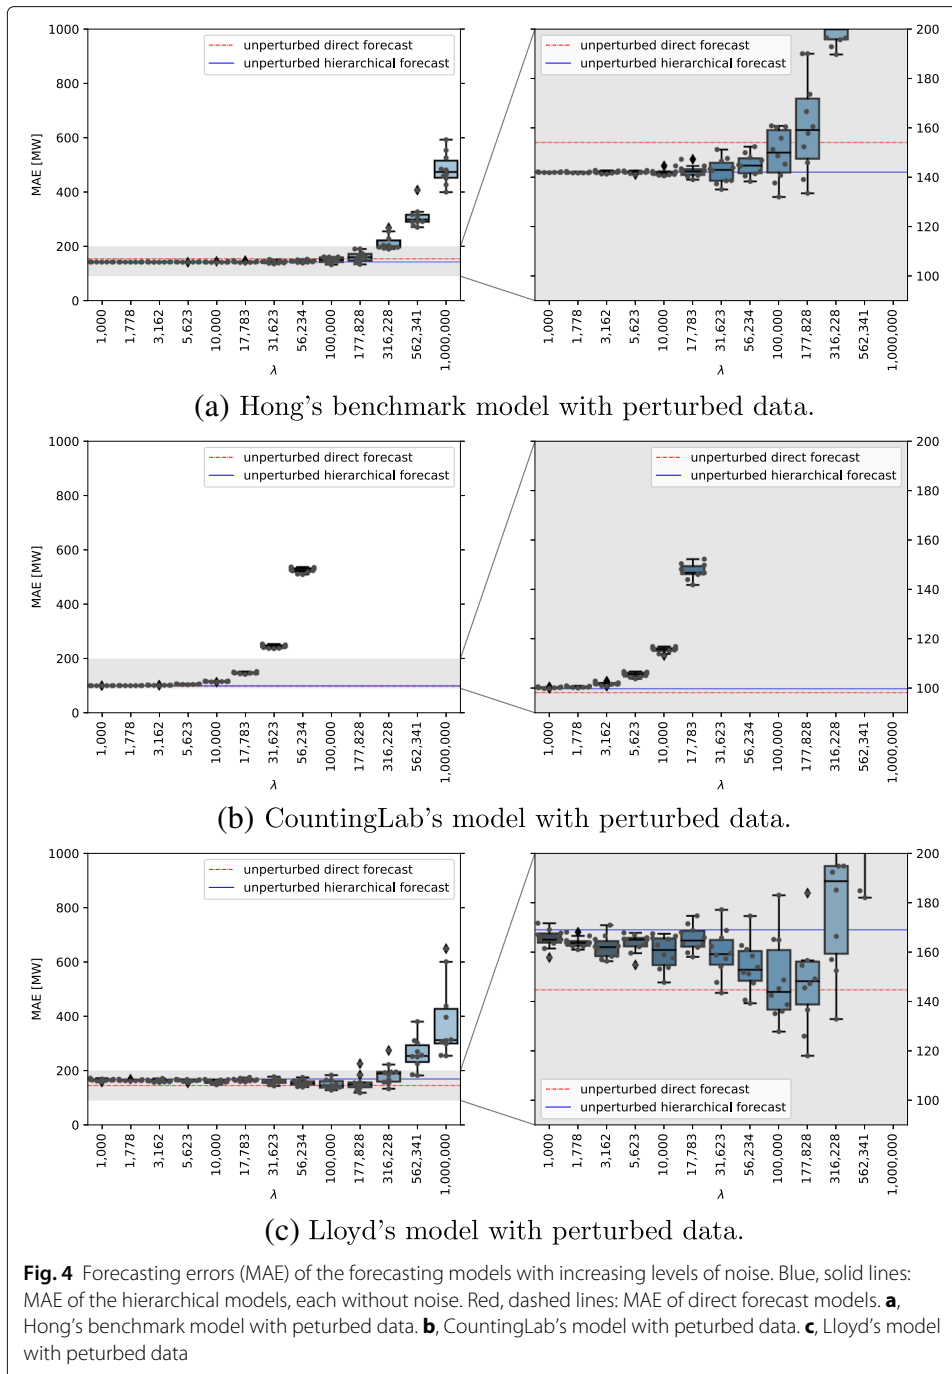
The MAPE and MAE of Lloyd's forecast with perturbed data for each zone are depicted in Fig. 4c. Surprisingly, the forecast first improves for some amount of noise, reaches a minimum at $\lambda = 177, 828$ and then rises quickly.

This behavior can be attributed to the gradient boosting model which also has the highest weight (0.765) in the ensemble averaging process (not shown). Since the inputs of the gradient boosting model do not include any load values (compare "Lloyd's forecasting method" section), Differential Privacy acts as output noise

Table 1 Noise (λ) and sensitivities (Δf) lead to ϵ and interpretable re-identification confidence (ρ) for k-fold adaptive composition ($k = 38,070$)

λ	Δf [kW]	ϵ	$\tilde{\epsilon}_{\delta}$	ρ	u^{λ}		
					CountingLab	Lloyd	Benchmark
0	-	-	-	-	-1.53	-16.81	7.80
10,000	7.57	0.00076	0.92	0.72	-17.62	-10.48	7.98
10,000	10.05	0.00100	1.23	0.77			
10,000	15.36	0.00154	1.92	0.87			
10,000	48.00	0.00480	6.46	1.00			
56,234	7.57	0.00013	0.15	0.54	-433.29	-6.49	5.94
56,234	10.05	0.00018	0.21	0.55			
56,234	15.36	0.00027	0.32	0.58			
56,234	48.00	0.00085	1.04	0.74			
100,000	7.57	0.00008	0.08	0.52	-1084.80	-2.76	3.10
100,000	10.05	0.00010	0.11	0.53			
100,000	15.36	0.00015	0.18	0.54			
100,000	48.00	0.00048	0.57	0.64			

The first row states utility (u) of the hierarchical, unperturbed forecast ($\lambda = 0$) over the direct forecast



which has been shown to potentially improve a model by (Breiman 2000). As the benchmark model did, the third classifier of the ensemble, the Gaussian Process model, degrades monotonically and finally rather quickly with increasing λ (not shown). The bad reaction upon noise of the Gaussian Process is plausible since the model heavily relies on a limited amount of 500 load values which corresponds to three weeks of data. However, since it only has a weight of 0.135 the gradient boosting model dominates for small λ .

Application of differential privacy

While we conceptually presented the integration of Differential Privacy into smart metering load forecasting in “Differentially private metering process” section, we provide an evaluation of the implementation in the following.

As initial step, we let an energy provider set utility bounds by choosing the noise scale λ in dependence of the acceptable loss in utility, i.e., forecast accuracy. In the next step we fix $\Delta f = 48$ kW as *global* Δf (i.e., maximum power consumption), which is the maximum power limit of 3-phased circuits in German residential homes. Based on λ and Δf , a global privacy guarantee of $(\epsilon, 0)$ -Differential Privacy (1) is provided by each individual load aggregate $\hat{L}_{z,t}$ using the Laplace mechanism (2).

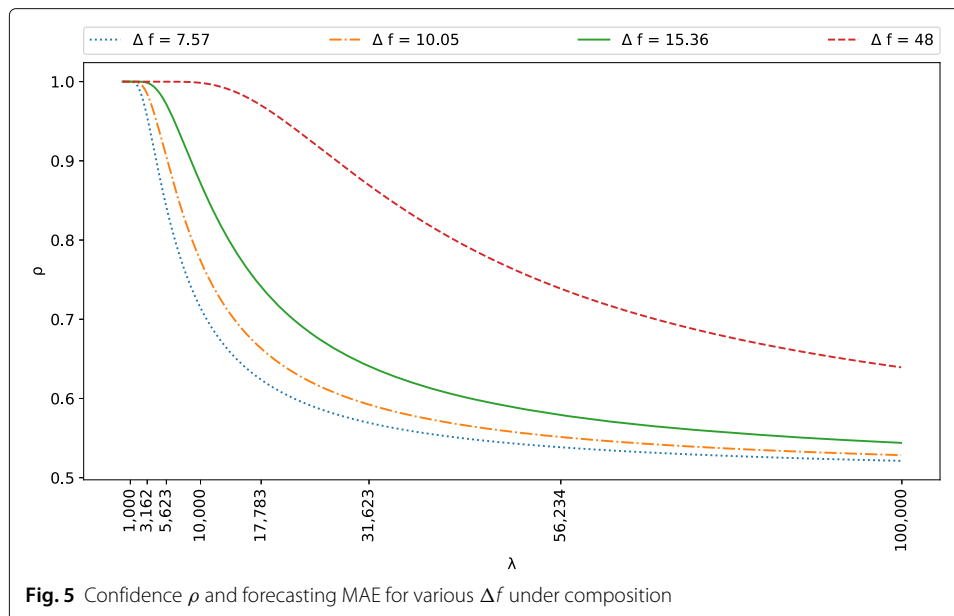
However, this theoretical restriction is far from being reached in practice. Thus, households may exchange the global Δf by a smaller, *local* Δf to identify their actual privacy guarantee. Considering the same λ , since $\epsilon = \Delta f/\lambda$, households may actually enjoy a stronger (smaller ϵ) protection against membership inference under their local Δf . However, the ϵ guarantee then only applies to loads within the local interval and does not keep an attacker from finding out about the bounds of that local interval. In the end, it is a matter of interpretation whether one relies on a very theoretical protection guarantee or a more realistic relaxation. To illustrate the impact, we vary Δf^5 according to Table 2 for our scenario.

When continuously releasing information by computing $\hat{L}_{z,t'_1}, \dots, \hat{L}_{z,t'_k}$ a composition theorem has to be applied as each $\hat{L}_{z,t'}$ relates to the same set of individuals (i.e., households). The GEFCom data set consists of $k = 38,070$ hourly load recordings, thus we have almost 40,000 composition steps. For large k , however, k -fold adaptive composition (“Differential Privacy” section) is a tight estimation of the privacy loss. By fixing some very small $\hat{\delta}$, the growth of a composed $\tilde{\epsilon}_{\hat{\delta}}$ no longer (3) depends linearly on k . We set $\tilde{\delta} \leq \frac{1}{|D|}$, where in the worst case w.r.t. the GEFCom data set $|D|$ is the number of all households in the US in 2013⁶, i.e., $\tilde{\delta} = \frac{1}{117,716,237} \approx 10^{-9}$. In the end, each household is protected by $(\tilde{\epsilon}_{\tilde{\delta}}, \tilde{\delta})$ -Differential Privacy.

Regarding our aim to express the privacy guarantee in a comprehensible way, $\tilde{\epsilon}_{\tilde{\delta}}$ is transformed into ρ by (5). The impact of λ on ρ is displayed in Fig. 5 for various Δf to illustrate the significant difference in membership inference likelihood when using theoretical worst case power consumption (i.e., $\Delta f = 48$ kW) or realistic maximum demands (i.e., $\Delta f = 15.36$ kW). Lowering Δf to more realistic values causes ρ to decrease and consequently results in stronger protection against membership inference. Thus, for $\lambda \geq 50,000$, households with realistically estimated maximum loads (Δf) have already acceptable privacy levels. At $\lambda = 100,000$, even the theoretical worst case of 48 kW approaches the desired $\rho = 0.6$ (cf. “Differentially private metering process” section).

Table 2 Selected Δf and according reasoning based on (Smart Metering Project - Electricity Customer Behaviour Trial 2012)

Δf	Argument
7.57 kW	90th percentile of highest power demands recorded per hour
10.05 kW	99th percentile of highest power demands recorded per hour
15.36 kW	highest power demand per hour in the whole data set
48.00 kW	maximum power demand fused in German residential homes



The trade-off between privacy and utility is shown in Table 1. Both CountingLab's and Lloyd's model work better for the direct than for the hierarchical setting. In contrast, the hierarchical benchmark forecast outperforms its direct counterpart. Thus, only the benchmark model is a suitable candidate for Differential Privacy. This is an interesting and unexpected result (note that although the performance of Lloyd's forecast improves with limited amount of noise it never has a positive utility). The desired membership inference confidence region $\rho \leq 0.6$ is achieved for the benchmark model for $\lambda = 56,234$ with $\Delta f = 15.35$ and offers a positive utility of 5.94% with respect to the direct forecast. Thus, a setting has been found where both, privacy and utility, have been reached. The authors want to highlight that they assume communication of individual, understandable membership inference risk ρ based on individual Δf as crucial to foster consumer acceptance of privacy-preserving techniques.

Related work

One of the first works to discuss and demonstrate privacy issues with smart metering was from (Molina-Markham et al. 2010). Most recently, (Rafsanjani et al. 2016) showed empirically that the occupancy of a commercial building can be estimated based on high-resolution energy consumption data with an accuracy above 95%.

Two prominent use cases of smart metering data are electricity consumption billing and real-time monitoring for grid operations. For billing, exact fees are important. Hence, due to the addition of noise, Differential Privacy has rarely been applied (Danezis et al. 2011). Typically, privacy is improved by disclosing only the necessary information for the business process, which is, at best, the final cost of each individual. Molina-Markham et al. (2010); Rial and Danezis (2011) and Jawurek et al. (2011) use Zero-Knowledge Protocols to provide privacy-preserving billing.

For real-time electricity monitoring, information aggregated over a geographical or topological grid area are sufficient. The privacy enhancing approaches for this use case are

based mostly on mixing networks which are partially backed by homomorphic encryption. Examples include work by Li et al. (2010); Garcia and Jacobs (2010); Defend and Kursawe (2013), and Finster and Baumgart (2014). Approaches related to aggregation are based on privacy definitions like k -anonymity. One representative is (Jia et al. 2017). These approaches do not provide real guarantees for privacy as they depend on knowledge limitations for the adversary.

All the approaches so far require the metering infrastructure to be designed in a specific way. As a privacy self-defence mechanism, a grid customer could resort to load obfuscation. Load obfuscation physically manipulates the load profiles of households by using battery storage systems or controllable loads and generators. Examples are Kalogridis et al. (2010) and McLaughlin et al. (2011), who leverage batteries to shift loads, (Chen et al. 2014), who controls Combined Heat and Power plants, and (Egarter et al. 2014), who uses energy management of appliances to protect privacy.

The closest related to our work are differentially private smart metering concepts. Ács and Castelluccia were the first to apply Differential Privacy on smart metering data. In their work (Ács and Castelluccia 2011), a distributed Laplace mechanism is applied using Gamma distributions before the data is mixed with other smart meters in an aggregation group. Bao and Lu (2015) investigated further the security and fault tolerance properties of the aggregation and mixing protocol. Eibl and Engel (2017) introduced post-processing to be applied on the perturbed data to improve the utility while still guaranteeing the same privacy level. They also discuss the required number of households in an aggregation group in order to be useful to the data analyst. Böhler et al. (2017) suggest using Differential Privacy with relaxed sensitivity and a privacy-preserving correction algorithm in IoT scenarios to still allow outlier detection while protecting the majority of households. Barbosa et al. (2016) also discussed filtering techniques to improve utility after the noise has been added to the aggregate. Their work evaluates the protection of individual appliances in single households by considering multiple device sensitivities in load profiles and by using Differential Identifiability. However, they do not address the compatibility condition $m = 2$ to allow utilizing Differential Identifiability in Differential Privacy scenarios. Besides Differential Identifiability, another method for rationally choosing ϵ was proposed in (Hsu et al. 2014). Yet, this approach is purely economically driven and introduces a handful of new parameters depending again on subjective assumptions on a given scenario. In contrast, focusing more on unconditional privacy, we further analyze Differential Identifiability. From Ács et al. (2011) we borrowed the way how to generate Laplacian noise in a distributed way. While in (Ács and Castelluccia 2011) the focus is on the aggregation protocol, we further improve composition and connect to Differential Identifiability and load forecasting with utility guarantees.

Conclusion and outlook

In this paper, we discussed that energy providers are interested in smart metering data to refine the forecast of domestic loads of their customers. As this conflicts with the privacy loss incurred by the acquisition of individual load profiles, we designed a differentially private metering process based on building blocks already proposed in previous works. Using three well-documented load forecasting approaches, we evaluate whether using

smart metering data provides an actual benefit for the energy provider. We found out that this is not always the case and that the forecasting approaches are variously susceptible to noise. If smart metering data actually provides a utility to the energy provider, Differential Privacy allows to gradually trade-off utility against forecasting performance. Our results show that for one forecasting approach, reasonable utility can be reached while providing a strong privacy guarantee. In that case, Differential Identifiability even provides an intuitive interpretation of the amount of privacy loss.

Several important points have to be considered when our concept is to be applied safely in practice: Firstly, there is no privacy guarantee for individual smart metering data of a *single* household. In particular, the sum of individual load and Gamma noise is still sensitive, therefore secure aggregation with other households is crucial. That is why we stated homomorphic encryption and masking or mixing as minimum requirement (cf. “Differentially private metering and load forecasting” section). Secondly, we considered privacy guarantees from a static snapshot of the scenario when the energy provider has collected approximately 4.5 years of zonal load profiles. Applying our approach in practice continuously would mean that the privacy guarantee is stronger if less than 4.5 years of data was collected from the customer. After 4.5 years our evaluated privacy guarantees would slowly degenerate. Thirdly and tightly connected to the second point, the historic and forecasted load profiles of our used data set were given with hourly readout intervals. However in Europe, load profiles are acquired in 15 min intervals. Our findings also apply to this case with the only difference that the privacy guarantee would hold for slightly more than one year instead of 4.5 years. Finally, if the privacy level offered by the energy provider is not high enough to protect the electricity usage of the whole household, the protection can still be interpreted for single household appliances. In this case, one has to be aware that the usage of this single appliance is not allowed to correlate to the (parallel) usage of other appliances.

There are several natural extensions to the presented work: Firstly, for utility evaluation, we used three well-documented point forecasting methods. Point forecasting outputs only a single (the likeliest) load value for one time interval. An extension to this work would be to evaluate differentially private metering with probabilistic forecasting methods (cf. (Hong and Fan 2016)). Secondly, our concept perturbs and transmits the complete zonal time series to the energy provider and the forecasting model training is performed by the energy provider. In the future, we plan to integrate Differential Privacy directly into a distributed model training approach on the customer side using *objective-function perturbation* for less privacy loss and tighter guarantees. Thirdly, lowering the local sensitivity by minimizing the household’s peak load leads to a stronger privacy level. Incidentally, automatic energy management systems like the ones described in (Egarter et al. 2014) and (Mauser et al. 2016) are able to shift controllable loads or control battery storages and combined heat and power plants to facilitate this idea. Fourthly, with the continual release of load data in practice, the privacy loss quantified by ϵ would slowly add up over the course of time. To be aware of one owns privacy situation, one needs to keep track of how much privacy was already leaked to which party. The data custodian proposed by (Rigoll and Schmeck 2017) provides such an accounting service. Finally, the perturbed data could be filtered (e.g., using moving average or Kalman filters) to compensate the noise as already proposed in (Bao and Lu 2015; Eibl and Engel 2017).

As final remark, using local sensitivities creates the incentive to limit one own's energy consumption due to privacy protection interests. Although this behavior may be beneficial to the electric grid, this would not be in the spirit of informational self-determination. That is why using the global sensitivity instead of local sensitivities should be preferred.

Endnotes

¹ The control energy price in Germany ("reBAP") is available for download on <https://www.regelleistung.net>

² e.g., household occupancy, appliance usage or approximate sleep-wake-cycles

³ Re-implementation available at <https://github.com/KaibinBao/differentially-private-stlf>

⁴ Source code available at <https://github.com/jamesrobertlloyd/GEFCOM2012>

⁵ Local Δf are based on statistics retrieved from the CER data set (Smart Metering Project - Electricity Customer Behaviour Trial 2012) as the GEFCOM data set only contains aggregates. For the Adaptive Composition, one has to regard the peak load (power) within a read-out interval. There are *technical* bounds on the instantaneous electric power usage we can use to derive global sensitivities for our privacy model. In Germany, a household is usually protected with a 63 A contractor and the household is connected to all three AC-phases (cf. Section 15.2 in (Kasikci 2013) and DIN 18015 Part 1). With a nominal voltage of 230 V and an acceptable over-voltage of 10%, we get that the highest electrical power consumption a German household could have is $P_{peak} = 230V \cdot 1.10 \cdot 63A \cdot 3 \approx 48 \text{ kW}$. The global sensitivity $\Delta f = 48 \text{ kW}$ for a read-out interval is very likely to be higher than actual peaks households' power demand. Consequently, we calculate individual households' risk with their corresponding local sensitivities in the privacy evaluation. This enables households to derive their individual Differential Privacy guarantee. To obtain these values, we unfortunately cannot use the GEFCOM data set due to missing information on load recordings of single households. Thus, we calculated the 90th, 99th and 100th percentile of the highest consumption peaks over all households from the comparable CER electric data set (Smart Metering Project - Electricity Customer Behaviour Trial 2012) to obtain good approximation of a realistic maximum.

⁶ Estimated 117,716,237 by the U.S. Census Bureau: <https://www.census.gov/quickfacts/fact/table/US/HSD410216>

Abbreviations

AC: Alternating current; CER: Commission for energy regulation – Ireland's independent energy and water regulator; DSO: Distribution system operator; GEFCOM: Global energy forecast competition; IoT: Internet of things; MAE: Mean absolute error; MAPE: Mean absolute percentage error; reBAP German: "regelzonenübergreifender einheitlicher Bilanzausgleichsenergiepreis" – Uniform balancing energy price across control zones; TSO: Transmission system operator

Funding

Publication costs for this article were sponsored by the Smart Energy Showcases - Digital Agenda for the Energy Transition (SINTEG) programme. This work has received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 653497 PANORAMIX, from the Federal Ministry of Education and Research (BMBF) under funding No. 16KIS0843 KASTEL-SKI, and from the federal state of Salzburg.

Availability of data and materials

The GEFCOM 2012 data set analyzed during the current study is available in Appendix A of the GEFCOM 2012 publication (Hong et al. 2014). The smart metering data set of the Commission for Energy Regulation (CER) from the Electricity Customer Behaviour Trial (Smart Metering Project - Electricity Customer Behaviour Trial 2012) analyzed during the current study is available in the Irish Social Science Data Archive (ISSDA), <http://www.ucd.ie/issda/data/commissionforenergyregulationcer>. The data set generated during the current study can be reproduced using the code provided, <https://github.com/KaibinBao/differentially-private-stlf>.

About this Supplement

This article has been published as part of *Energy Informatics* Volume 1 Supplement 1, 2018: Proceedings of the 7th DACH+ Conference on Energy Informatics. The full contents of the supplement are available online at <https://energyinformatics.springeropen.com/articles/supplements/volume-1-supplement-1>.

Authors' contributions

The concept and methodology of this specific case study of differentially private smart metering data emerged from meetings of GE, KB, PG, and DB. Additional individually attributable contributions are: GE provided the formal problem definition. He also implemented, conducted and described the experiments of CountingLab's and Lloyd's forecasting method. KB drafted the context of the case study in the manuscript and analyzed the sensitivity of single households. He implemented, conducted and described the experiments of Hong's benchmark method. DB and PG provided the interpretation of the privacy guarantee in terms of Differential Identifiability and drafted the interpretation of the Differential Privacy guarantees in the manuscript. They also investigated k-fold adaptive composition for a tighter lower bound of the privacy guarantee. HS helped to write the final version of this publication. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Salzburg University of Applied Sciences, Center for Secure Energy Informatics, Urstein Süd 1, Puch/Salzburg, Austria.

²Karlsruhe Institute of Technology (KIT), Institute AIFB, Kaiserstr. 12, 76131 Karlsruhe, Germany. ³SAP Security Research, Vincenz-Prießnitz-Str. 1, 76131 Karlsruhe, Germany.

Published: 10 October 2018

References

- Acs G, Castelluccia C (2011) I have a DREAM! (Differentially privatE smArt Metering). In: Proc. of the 13th International Conference on Information Hiding (IH). Springer, Berlin, pp 118–132
- Bundesnetzagentur (2011) Anlage zum Beschluss BK6-07-002: Marktregeln für die Durchführung der Bilanzkreisabrechnung Strom (MaBiS) in der konsolidierten Lesefassung vom 28.10.2011. https://www.bundesnetzagentur.de/DE/Service-Funktionen/Beschlusskammern/Beschlusskammer6/BK6_31_GPKE_und_GeLiGas/Mitteilung_Nr_31/Anlagen/Konsolidierte_Lesefassung_MaBiS.pdf?__blob=publicationFile&v=2. Accessed 23 May 2018
- Breiman L (2000) Randomizing outputs to increase prediction accuracy. *Mach Learn* 40(3):229–242
- Bao H, Lu R (2015) A New Differentially Private Data Aggregation with Fault Tolerance for Smart Grid Communications. *IEEE Internet of Things J* 2(3):248–258
- Buescher N, Boukoros S, Bauregger S, Katzenbeisser S (2017) Two Is Not Enough: Privacy Assessment of Aggregation Schemes in Smart Metering. *Proc Priv Enhancing Technol* 2017(4):118–134
- Böhler J, Bernau D, Kerschbaum F (2017) Privacy-preserving outlier detection for data streams. In: IFIP Annual Conference on Data and Applications Security and Privacy. Springer, Cham, pp 225–238
- Barbosa P, Brito A, Almeida H (2016) A Technique to provide differential privacy for appliance usage in smart metering. *J Inf Sci* 370(Supplement C):355–367
- Chen D, Irwin D, Shenoy P, Albrecht J (2014) Combined heat and privacy: Preventing occupancy detection from smart meters. In: Proc. of the 2014 IEEE International Conference on Pervasive Computing and Communications (PerCom). IEEE, New Jersey, pp 208–215
- Charlton N, Singleton C (2014) A refined parametric model for short term load forecasting. *Int J Forecast* 30(2):364–368
- Defend B, Kursawe K (2013) Implementation of privacy-friendly aggregation for the smart grid. In: Proc. of the 1st ACM Workshop on Smart Energy Grid Security (SEGS). ACM, New York, pp 65–74
- Dwork C (2006) Differential Privacy. In: Proc. of the 33rd International Colloquium on Automata, Languages and Programming (ICALP). Springer, Berlin, pp 1–12
- Dwork C, McSherry F, Nissim K, Smith A (2006) Calibrating noise to sensitivity in private data analysis. In: Proc. of the 3rd Conference on Theory of Cryptography (TCC). Springer, Berlin, pp 265–284
- Dwork C, Naor M, Pitassi T, Rothblum GN (2010) Differential privacy under continual observation. In: Proc. of the 42nd ACM Symposium on Theory of Computing (STOC). ACM, New York
- Danezis G, Kohlweiss M, Rial A (2011) Lecture Notes in Computer Science (including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). In: Filler T, Pevny T, Craver S, Ker A (eds). Springer, Berlin, pp 148–162
- Dwork C, Roth A (2014) The Algorithmic Foundations of Differential Privacy. *Found Trends Theor Comput Sci* 9(2013):211–407
- Commission Regulation (2017) (EU) 2017/1485 of 2 August 2017 establishing a guideline on electricity transmission system operation. Official Journal of the European Union. <http://data.europa.eu/eli/reg/2017/1485/oj>. Accessed 23 May 2018
- Commission Regulation (2017) (EU) 2017/2195 of 23 November 2017 establishing a guideline on electricity balancing. Official Journal of the European Union. <http://data.europa.eu/eli/reg/2017/2195/oj>. Accessed 23 May 2018
- Egarter D, Prokop C, Elmenreich W (2014) Load hiding of household's power demand. In: Proc. of the 5th IEEE International Conference on Smart Grid Communications. IEEE, New Jersey, pp 854–859
- Eibl G, Engel D (2017) Differential Privacy for Real Smart Metering Data. *Comput Sci Res Dev* 32(1-2):173–182
- Erkin Z, Tsudik G (2012) Private Computation of Spatial and Temporal Power Consumption with Smart Meters. In: Proc. of the 10th International Conference on Applied Cryptography and Network Security (ACNS). Springer, Berlin, pp 561–577
- Eibl G, Engel D (2015) Influence of Data Granularity on Smart Meter Privacy. *IEEE Trans Smart Grid* 6(2):930–939

- Finster S, Baumgart I (2014) SMART-ER: Peer-based privacy for smart metering. In: Proc. of the 2014 IEEE Conf. on Computer Communications Workshops (INFOCOM). IEEE, New Jersey. pp 652–657
- Friedman JH (2001) Greedy function approximation: A gradient boosting machine. *Ann Stat* 29(5):1189–1232
- Fan S, Methaprayoon K, Lee W-J (2009) Multiregion load forecasting for system with large geographical area. *IEEE Trans Ind Appl* 45(4):1452–1459
- Federal Energy Regulatory Commission (2015) Energy primer, a handbook of energy market basics. Federal Energy Regulatory Commission, Washington, DC. pp 1–140
- Garcia FD, Jacobs B (2010) Privacy-friendly energy-metering via homomorphic encryption. In: Proc. of the 6th Conference on Security and Trust Management (STM). Springer, Berlin. pp 226–238
- Hong T, Fan S (2016) Probabilistic electric load forecasting: A tutorial review. *Int J Forecast* 32(3):914–938
- Hsu J, Gaboardi M, Haeberlen A, Khanna S, Narayan A, Pierce BC, Roth A (2014) Differential privacy: An economic method for choosing epsilon. In: Proceedings of the 27th IEEE Computer Security Foundations Symposium (CSF). IEEE, New Jersey. pp 398–410
- Hong T, Pinson P, Fan S (2014) Global energy forecasting competition 2012. *Int J Forecast* 30(2):357–363
- Ilić D, da Silva PG, Karnouskos S, Jacobi M (2013) Impact assessment of smart meter grouping on the accuracy of forecasting algorithms. In: Proc. of the 28th ACM Symposium on Applied Computing (SAC). ACM, New York. pp 673–679
- Jawurek M, Johns M, Kerschbaum F (2011) Plug-in privacy for smart metering billing. In: Privacy Enhancing Technologies Symposium. Springer, Berlin. pp 192–210
- Jia R, Sangogboye FC, Hong T, Spanos C, Kjærgaard M. B. (2017) PAD: Protecting Anonymity in Publishing Building Related Datasets. In: Proc. of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments. BuildSys '17. ACM, New York. pp 4–1410
- Kasikci I (2013) Planung Von Elektroanlagen. Springer, Wiesbaden
- Kalogridis G, Efthymiou C, Denic SZ, Lewis TA, Cepeda R (2010) Privacy for smart meters: Towards undetectable appliance load signatures. In: Proc. of the 1st IEEE International Conference on Smart Grid Communications. IEEE, New Jersey. pp 232–237
- Kotz S, Kozubowski T, Podgorski K (2001) The Laplace Distribution and Generalizations. Birkhäuser, Basel
- Knirsch F, Eibl G, Engel D (2018) Error-resilient masking approaches for privacy preserving data aggregation. *IEEE Trans Smart Grid* 9(4):3351–3361
- Kairouz P, Oh S, Viswanath P (2017) The Composition Theorem for Differential Privacy. *IEEE Trans Inf Theory* 63(6):4037–4049
- Lisovich MA, Mulligan DK, Wicker SB (2010) Inferring personal information from demand-response systems. *IEEE Secur Priv* 8(1):11–20
- Li F, Luo B, Liu P (2010) Secure Information Aggregation for Smart Grids Using Homomorphic Encryption. In: Proc. of the 1st IEEE International Conference on Smart Grid Communications. IEEE, NJ. pp 327–332
- Lee J, Clifton C (2012) Differential Identifiability. In: Proc. of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD). ACM, New York. pp 1041–1049
- Li N, Qardaji W, Su D, Wu Y, Yang W (2013) Membership privacy: A unifying framework for privacy definitions. In: Proc. of the 2013 ACM SIGSAC Conference on Computer & Communications Security (CCS). ACM, New York. pp 889–900
- Lloyd JR (2014) GEFCom2012 hierarchical load forecasting: Gradient boosting machines and Gaussian processes. *Int J Forecast* 30(2):369–374
- McLaughlin S, McDaniel P, Aiello W (2011) Protecting consumer privacy from electric load monitoring. In: Proceedings of the 18th ACM conference on Computer and communications security. ACM, New York. pp 87–98
- McDaniel P, McLaughlin S (2009) Security and Privacy Challenges in the Smart Grid. *IEEE Secur Priv* 7(3):75–77
- Molina-Markham A, Shenoy P, Fu K, Cecchet E, Irwin D (2010) Private memoirs of a smart meter. In: Proc. of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-efficiency in Building (BuildSys). ACM, New York. pp 61–66
- MsbG (2016) Messstellenbetriebsgesetz vom 29. August 2016 (BGBl. I S. 2034), das durch Artikel 15 des Gesetzes vom 22. Dezember (BGBl. I S. 3106) geändert worden ist. <https://www.gesetze-im-internet.de/messbg>. Accessed 23 May 2018
- Mausser I, Müller J, Allerding F, Schmeck H (2016) Adaptive building energy management with multiple commodities and flexible evolutionary optimization. *Renew Energy* 87:911–921
- Rigoll F, Schmeck H (2017) A concept for a user-oriented energy data management system. In: Helmholtz Portfolio Theme Large-scale Data Management and Analysis. KIT Scientific Publishing, Karlsruhe. pp 23–39
- Rial A, Danezis G (2011) Privacy-preserving smart metering. In: Proc. of the 10th Annual ACM Workshop on Privacy in the Electronic Society. ACM, New York. pp 49–60
- Rafsanjani HN, Ahn CR, Chen J (2016) Linking building energy-load variations with occupants' energy-use behaviors in commercial buildings: non-intrusive occupant load monitoring (NIOLM). *Energy Build* 172:317–327
- Smart Metering Project - Electricity Customer Behaviour Trial CER (2012) 2009-2010 dataset. Commission for Energy Regulation (CER). Irish Social Science Data Archive. <https://www.ucd.ie/issda/data/commissionforenergyregulationcer>. Accessed 11 Jan 2018
- StromNZV (2017) Stromnetzzugangsverordnung vom 25. Juli 2005 (BGBl. I S. 2243), die zuletzt durch Artikel 1 der Verordnung vom 19. Dezember (BGBl. I S. 3988) geändert worden ist. <https://www.gesetze-im-internet.de/stromnzzv>. Accessed 23 May 2018
- Epe C, Fuhrberg-Baumann J, Herbst U, Hermann M, Kreye HD, Mahn U, Mönnig R, Scherer U (2007) DistributionCode. Verband der Netzbetreiber VDN e.V. beim VDEW, Berlin. pp 1–28