Smart Meter Data Analytics for Enhanced Energy Efficiency in the Residential Sector

Mariya Sodenkamp[1], Ilya Kozlovskiy[1], Konstantin Hopf[1], Thorsten Staake[12]

[1]Energy Efficient Systems Group, University of Bamberg
[2]Department of Management, Technology and Economics, ETH Zurich
{mariya.sodenkamp, ilya.kozlovskiy, konstantin.hopf, thorsten.staake}@uni-bamberg.de

Abstract. Achievement of the ambitious environmental sustainability targets requires improvement of energy efficiency practices in private households. We demonstrate how utility companies, having access to smart electricity meter data, can automatically extract household characteristics related to energy efficiency and adoption of renewable energy technologies (e.g., water/space heating type, age of house, number and age of electric appliances, interest in installation of photovoltaic systems etc.) by using supervised-machine-learning-based green IT artifacts. The gained information enables design of custom-tailored interventions (such as promotion of personalized energy audits, ecologic services and products, or load shifting mechanisms) that trigger residents' behavioral change toward environmental sustainability as well as improvement of utilities' key performance indicators. Moreover, realizing privacy preservation concerns, we investigate the influence of smart meter data granularity and the amount of survey responses required for the artifact development on the household classification quality.

Keywords: Green information systems (IS), smart meters, data analytics, energy efficiency, sustainability

13th International Conference on Wirtschaftsinformatik,
February 12-15, 2017, St. Gallen, Switzerland

Sodenkamp, M.; Kozlovskiy, I.; Hopf, K.; Staake, T. (2017): Smart Meter Data Analytics for Enhanced Energy Efficiency in the Residential Sector, in Leimeister, J.M.; Brenner, W. (Hrsg.): Proceedings der 13. Internationalen Tagung Wirtschaftsinformatik (WI 2017), St. Gallen, S. 1235-1249

1 Introduction

In order to promote energy efficiency and the integration of renewable energy sources, policy makers have placed high hopes on networked electricity meters that measure and communicate consumption information at a high resolution in time [1]. These so-called smart meters offer, besides improvements to the utilities' billing processes, timely consumption feedback for residential customers, render dynamic tariff schemas possible, and provide input to home automation systems. Ultimately, smart meters should help citizens to reduce their electricity consumption and motivate shifting loads to support the integration of renewable, fluctuating electricity sources. Early studies have reported promising effects including energy savings as high as 15% [2], and many countries mandated utility companies to roll out the technology among their residential customers. To date, more than 50 million smart meters have been deployed in the EU, and industry expects 154 million devices by 2017 [3].

The initial enthusiasm for the technology, however, has faded recently. More carefully designed studies that were not affected by sampling bias (e.g., from primarily observing volunteering, motivated users) showed savings of only around 3% and little response to incentives for load shifting [4]. Moreover, the mostly poor design of feedback campaigns together with deficient data protection practices raised substantial privacy concerns among consumers [5]. It has become evident that relying on the mere effects of consumption information and dynamic tariffs is not sufficient to advance energy literacy, motivate the hoped-for behavioral change, nor to trigger investments in saving technologies [6].

We argue that the current, disappointing performance of smart metering is not a problem of the technology as such, but resides in an insufficient information extraction from the available metering data. Feedback interventions, for example, that are tailored to individual recipients have been shown to achieve substantially higher saving effects and to reach a better user acceptance [7]. Such measures (e.g., providing concrete advice or comparisons to similar households, offering services that reflect the household's characteristics, etc.) typically yield considerably lower cost per kWh saved (or shifted) than tax credits and rebates and meet higher public acceptance than prohibitive regulations [8]. Load shifting measures, control of heating systems or combined micro heat and power plants can also considerably benefit from information beyond pure metering data if, for example, up-to-date knowledge on the presence of the inhabitants is available. While strong evidence supports the benefits of using such specific information on households to conduct consumer-specific energy efficiency campaigns [8], a major problem is, that the required data to conduct such campaigns is not available for large scale deployments. This reduces the benefits of smart metering infrastructures dramatically.

The research outlined herein lies within the scope of energy informatics - an emerging discipline concerned with analysis, design and implementation of information systems to reduce energy consumption [9, 10]. ***We strive for providing the missing link between smart meter data and powerful energy efficiency measures***. We show how data collected at 15-minute granularity from out-of-the-shelf smart electricity meters can be used to infer energy-efficiency related characteristics of residential dwellings (e.g., water/space heating type, age of house, number and age of electric appliances, interest on installation of photovoltaic systems etc.), ***using supervised machine learning*** techniques. Automatically mining this information – with the consent of inhabitants – enables large-scale, targeted saving advice and better input to home automation systems that can significantly contribute to reduction of energy consumption and advancement of environmental sustainability.

2 Research Objective and Theoretical Context

A number of researchers have developed methods to predict household characteristics data from energy consumption data. The approaches differ with respect to the type of data available (e.g., load records, household survey reports), their resolution in time, and the output variables of interest. The potential to recognize characteristics from load curves depends on the data granularity – the finer the granularity is the better recognition performance can be achieved. Vast research has been done on the recognition of devices from the data of extreme granularity (Hertz and Megahertz frequency) – in the field of Non Intrusive Load Monitoring (NILM) [11, 12]. While sampling rates far beyond one megahertz are quite common in industrial or lab settings, the smart metering infrastructure that is currently being deployed and is expected to be in the field for the next 20 years does not provide such fine grained data. Therefore, the NILM methods will not be compatible with the standard meters deployed in most households [13].

Several authors have investigated coarser consumption data. Chicco [14] provides an overview of the clustering methods for electrical load pattern grouping. In the field of recognition of energy efficiency characteristics which we focus on, Fei et al.[15], proposed a method to detect heat pumps from the daily energy consumption data. Beckel et al. [16] used 30-minute data reduced to 26 features to infer 18 household properties, most of which relate to the inhabitants' life situation (age, family, employment, social class, etc.), and three directly relate to energy efficiency (number of appliances, cooking type and lightbulbs). Hopf et al. [17] and Sodenkamp et al. [18] improved Beckel's algorithm by extracting 88 features, applying filtering methods and by refining the properties. Sodenkmp et al. [19] included weather data to Beckel's algorithm using a multi-dimensional classification method called DID-class.

Further works employ conventional yearly consumption readings for dwelling classification. Kozlovskiy et al. [20] detected old gas heating systems for a targeted cross-selling campaign in Belgium. Sodenkamp et al. [21] predicted household probability to register on an energy efficiency portal for customer engagement campaigns in Germany and Switzerland. Hopf et al. [22] used yearly consumption and geographic data from OpenStreetMap and GeoNames to detect living area, household type and number of residents.

Our overarching goal is to **examine the usage of machine-learning-based smart meter data analytics methods in the practice of energy utility companies and their contribution to the environmental and economic sustainability.** In this work, we go beyond the state-of-the art by identifying eleven energy-efficiency related household characteristics (space heating type and age, water heating type, heat pump usage in a household, house age, number and age of appliances, presence and interest in photovoltaic and thermal installations, number of recently completed energy efficiency measures, and type of cooking facility) from smart meter data at 15-minute granularity, which we collected from a Swiss utility company. We extract 93 features from the consumption data and include weather data represented by 40 features. Finally, we investigate the effects of data granularity. This information can be used for the development of targeted energy-efficiency measures (e.g., saving tips, promotions of installations of renewable energy systems, load shifting campaigns). To enable algorithm training and testing we designed and conducted an online survey. Thus, our first research question is as follows:

RQ 1. *To what extent is it possible to recognize energy efficiency related household characteristics from smart meter data using machine learning methods?*

The availability of detailed electricity consumption data-traces to the utilities is associated with consumers' privacy concerns [5]. Depending on the company policy and local legislation, utilities make use of different data granularities. The typical aggregations vary between 15-min, 30-min, hourly and daily levels. For daily data, some utilities differentiate between the HT (high tariff, during the day) and NT (low tariff, during the night) consumption. Data with lower frequency contains less information about customers' behavior and we expect the performance of our classification algorithms to degrade when applied to such data. Therefore, we test how different data granularities influence the resulting classification performance and formulate our second research question as follows:

RQ 2*. To what extent do different granularities of smart meter data influence the recognition quality of energy efficiency related household characteristics?*

3 Data Description

For our data science study, we cooperated with a utility company in Switzerland with about 9`000 customers that provided us with household electricity meter readings at 15-minute granularity in the timespan between June 1st 2014 and May 31st 2015. For the same time period, we acquired hourly weather-data from the U.S. National Climate Data Center [23]. Together with the utility, we conducted a web-based customer survey about energy efficiency related household characteristics between June and September 2015. All customers received an invitation to the survey attached to their bimonthly bill. In this survey, we collected data on 527 households, which corresponds to a response rate of 6%. We matched the survey results with smart meter data using respondents' names and addresses.

Based on the survey responses, we defined and extracted 11 energy efficiency relevant household characteristics (*properties*) that include at least two *classes* (see Table 1). The class definition was either naturally given (e.g., heat pump exists / does not exist) or we empirically set class borders by using quantiles. When we used quantiles for defining the class border, we either aimed to separate the households in equally sized classes, or wanted to identify a specially interesting class (e.g., high purchasing intention for solar installation). The total number of households per property does not necessarily respond to the total number of households participating in the survey, since the survey participants left some questions unanswered. The amount of excluded instances due to the missing data is different per property and lies mostly at 23.48 %, except for "Age of residency" with 23.91%, "Age of appliances" with 34.35%, and "Age of heating" with 40% of all survey participants excluded.

Table 1: Energy efficiency related household properties with classes defined from survey responses; q_x denotes statistical x percentile of the survey responses (relative class sizes are not necessarily identical to the percentiles by definition for numerical values, due to categorical survey variables).

Property	Classes	Definition	Class size	
			Abs.	**Rel.**
Age of appliances	New	Avg. appliance age < $q_{0,25}$	153	33,41%
	Average	Avg. appliance age between $q_{0,25}$ and $q_{0,75}$	152	33,19%
	Old	Avg. appliance age > $q_{0,75}$	153	33,41%
Num. of appliances	Few	Number appliances < $q_{0,25}$	149	28,27%
	Average	Number appliances between $q_{0,25}$ and $q_{0,75}$	280	53,13%
	Many	Number appliances > $q_{0,75}$	98	18,60%
Cooking type	Electric	Number electric stoves > 0	484	91,84%
	Not electric	Number electric stoves = 0	43	8,16%
Efficiency measure	No	Number completed energy efficiency measures during the last 15 years (insulation of basement / roof / building envelop, or window replacement)	304	57,69%
	Few		109	20,68%
	Multiple		114	21,63%
Heat pump	No	Existing heat pump	453	85,96%
	Yes		74	14,04%
Age of residency	< 10	Age (in years) of the building the household is living in	71	13,65%
	10-29		147	28,27%
	30-74		219	42,12%
	≥ 75		83	15,96%
Interest in solar	Low	Purchase intention coefficient < $q_{0,50}$	387	73,43%
	Average	Purchase intention coefficient between $q_{0,50}$ and $q_{0,75}$	49	9,30%
	High	Purchase intention coefficient > $q_{0,75}$	91	17,27%
Solar installation	Yes	Photovoltaics or solar heating existent	29	5,50%
	No	Neither photovoltaics nor solar heating existent	498	94,50%
Age of heating	New	Space heating age < $q_{1/3}$	128	33,51%
	Average	Space heating age between $q_{1/3}$ and $q_{2/3}$	135	35,34%
	Old	Space heating age > $q_{2/3}$	119	31,15%
Space heating type	Electric	Space heating = „*Electric heating*"	21	3,98%
	Heat pump	Space heating = „*Heat pump*"	66	12,52%
	Other	Other space heating	440	83,49%
Water heating type	Electric	Water heating = „*Electric heating*"	81	15,37%
	Heat pump	Water heating = „*Heat pump*"	63	11,95%
	Other	Other water heating	383	72,68%

4 Analysis

In this section, we describe our data analysis methodology and results. We first seek to answer the first research question:

RQ 1. *To what extent is it possible to recognize energy efficiency related household characteristics from smart meter data using machine learning methods?*

To answer this question, we follow the four-step procedure described below.

Step 1: Feature Extraction. At the beginning, it is important to reduce the raw data dimensionality and transform the data to a more usable form. Consumption time series are divided into single weeks, since from previous work we know that a weekly energy consumption is sufficient to perform household classification [16]. We extracted *93 features from 15-min smart meter data* for each week, that are adopted from previous works dealing with 30-min smart meter data [16, 17]. The details can be found in the project report [24]. The extracted features cover four categories: consumption (e.g., in the morning, noon, evening); ratios of consumption figures (e.g., consumption in the morning vs. noon, daytime vs. night); statistics (e.g., variance, quantiles); others (e.g., number and average heights of consumption peaks).

Besides the smart meter features, we defined and used *40 features describing the correlation between electricity consumption and weather data*, since a positive effect of weather data on the classification performance was shown in the previous study [19]. For each weather variable (temperature, wind speed, sky cover, and precipitation), we calculate eight features: overall correlation over the week, correlation during the day and during times of the day (night, daytime, evening), correlation of minima in both time series, correlation of weather minima and consumption maxima, and ratio of the weekday and weekend correlations.

Finally, we use 133 features for our analysis. Due to the space constraints we cannot present all the features (interested reader is referred to [24]), but we list 10 most frequently selected features in the final prediction models for all household properties with a short description in Table 2.

Step 2: Feature selection. After having prepared the feature vectors, we select relevant features for each property separately. This is done to reduce overfitting and speed up the calculations by removing the irrelevant features. We tested the following feature selection methods: Correlation based (cfs), consistency based (consistency), Based on the importance from random forest (importance), statistical test for the difference in distributions (chi.squared), entropy based (gain.ratio), forward feature selection search (forward-selection), backward feature selection search (backward-selection), no feature selection (none). For the methods description see [25].

Table 2: The top 10 selected features and their description

Rank	Feature	Description
1	t15_above_2kw	Time with consumption above 2 kW
2	t15_value_min_guess	Time with consumption above minimal consumption
3	r15_wd_evening_noon	Relation between evening and noon consumption on weekdays
4	r15_mean_max_no_min	Relation between mean and max consumption with subtracted

		minimum
5	t15_time_above_base2	Time with consumption above estimated baseline
6	r15_evening_noon	Relation between evening and noon consumption
7	t15_const_time	Time with nearly constant consumption
8	t15_daily_min	Average daily minimum
9	r15_min_wd_we	Relation between minimum of consumption on weekdays and weekends
10	r15_var_wd_we	Relation between the variance of consumption on weekdays and weekends

Step 3: Classification. As the next step we train the model based on the selected features and evaluate the following six well-known classifiers that have implementations in the statistical programming environment GNU R: AdaBoost [26], k Nearest Neighbors (kNN) [27], Naïve Bayes [27], Random Forest [28], Support Vector Machine (SVM) [29].

Step 4: Evaluation. To measure the classification performance, we are interested in testing how many households are correctly classified in each available class. These numbers are typically presented in the form of a confusion matrix. To compare the different confusion matrixes between each other we calculate two measures from the confusion matrices:

- Accuracy**:** It is defined as the portion of correctly classified instances from the number of total classification instances and can take values between 0 and 1, where 1 corresponds to perfect prediction and 0 to total misclassification. Accuracy is easy to interpret, but in the situation where the classes are unbalanced (i.e., one class occurs much more often than the others) a classifier that always predicts a majority class can achieve high accuracy. Therefore, this measure can be slightly misleading if applied to such unbalanced properties.
- Matthews Correlation Coefficient (MCC): It is an alternative measure that is more suitable for the unbalanced problems. It is a correlation coefficient between the observed and predicted classifications. In the case of binary classification problem, it is equal with the phi statistic [30]. We use MCC definition for multiclass problems [31, 32]. MCC can take values between -1 and 1, where 1 corresponds to the perfect classification, -1 to the total disagreement between the predictions and real observations and 0 for the classification that is not better than random prediction. MCC lacks the easy interpretability of the accuracy measure, but it is a good compromise among discriminancy, consistency and coherent behaviors with varying number of classes, unbalanced datasets and randomization [31].

To calculate the performance measures, we first split the data into the *test set* (10%) from the *main data* by using a stratified split (the distribution of classes in the test set deviates at most with one household from the distribution in the main data). For the main data, we use 4-fold cross-validation to select the best classification algorithm and feature selection method.

Then, we take the features describing a single week (week number 34, from 12.01 to 18.01.2015). The features of the main dataset are centered and scaled to have mean

0 and standard deviation 1. The features in the test set are then scaled with the same proportions. All households with missing or not available values for the chosen week are removed, as well as the feature vectors that are constant for all households. This problem occurs mainly with weather features that are constant during the week (e.g., precipitation and sky cover).

Performance of the individual classifiers for week 34 is presented in Figure 1, together with two benchmark measures: random guess ($RG = 1/K$, where K is the number of classes in one property) and biased random guess ($BRG = \sum h^2$, where h is the relative class size of each class for one property, as displayed in Table 1). The figure illustrates, that no single classifier provides the best classification performance for all properties. Similarly, no best feature selection technique for all properties can be found. Therefore, we choose the classification configuration (feature selection and classifier) that produced the best result for MCC and list it in Table 3.

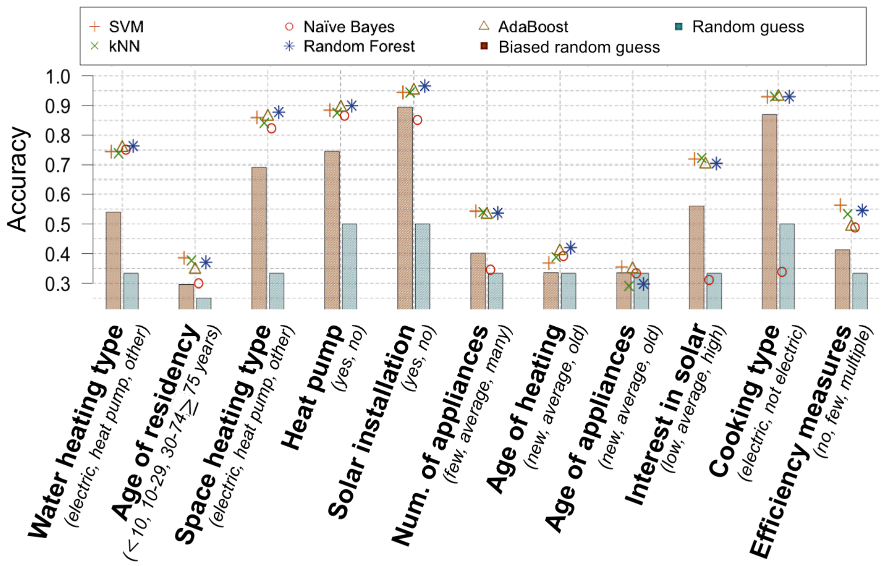

Figure 1: Comparison of classification performance of individual classifiers without feature selection (week number 34)

With these best performing classifier configurations, the final model is trained on all households that are not in the test set. Predictions are then made for the test set and are then evaluated in terms of both the accuracy and MCC. We repeated the training and prediction process for all 29 week of data that were not excluded due to public or school holidays or missing values. The results for every week are then aggregated with a simple ensemble classifier [33], that averages the prediction probabilities for each class and thus forms one single multi-week-classifier. The final predictions for the test data are then evaluated with respect to both the accuracy and MCC. The results are presented in the Figures 2 and 3.

As the result, we can answer our first research question positively: It is possible to predict most (9 from 11) energy efficiency relevant household properties better than by using random guess. We especially achieved good results for the prediction of properties related to the room and water heating (*Space heating type*, *Water heating type*, and *Heat pump*). The properties *Cooking type* and *Efficiency measures* could not be predicted adequately with our approach. But even with the negative MCC, the results can still be valuable for utility companies that usually do not know what household belong to which class: the application of the classifier can perform better than random guessing.

Table 3: The best performing configurations

Property	Feature selection method	Classifier
Space heating type	cfs	Random Forest
Water heating type	none	Random Forest
Heating age	gain.ratio	Support Vector Machine
Age of residency	cfs	AdaBoost
Age of appliances	chi.squared	Naïve Bayes
Cooking type	none	Random Forest
Heat pump	chi.squared	Random Forest
Solar installation	none	Random Forest
Efficiency measures	cfs	AdaBoost
Num. of appliances	none	Random Forest
Interest in solar	chi.squared	Random Forest

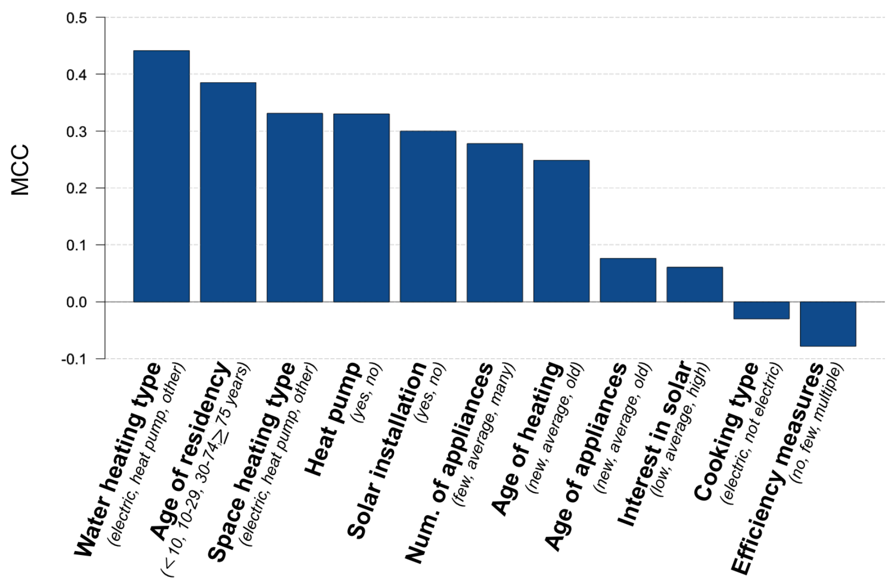

Figure 2: Classification performance for all properties measured with MCC

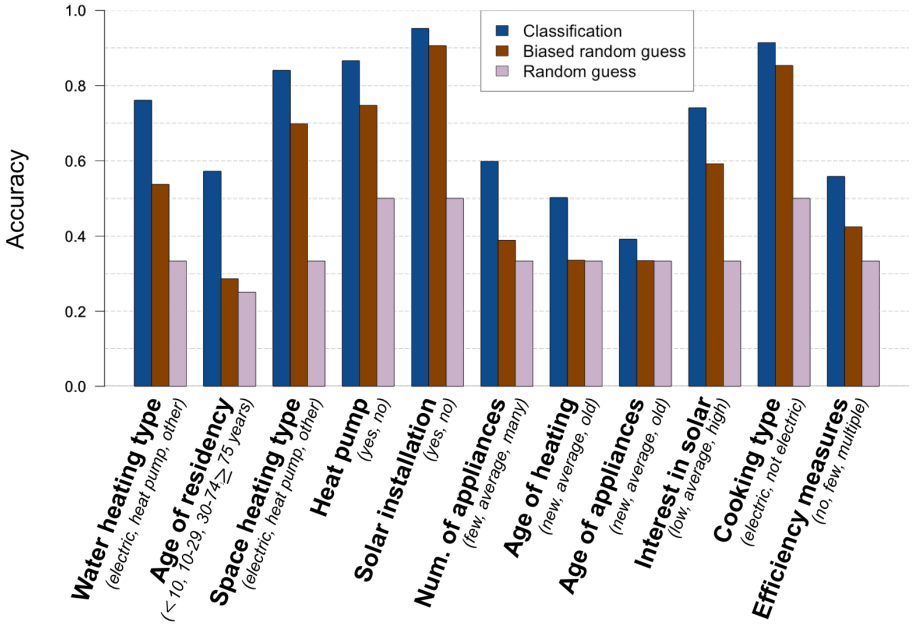

Figure 3: Classification results for all properties measured with accuracy, compared to random guess and biased random guess

RQ 2. *To what extent do different granularities of smart meter data influence the recognition quality of energy efficiency related household characteristics?*

To answer this question, we have to repeat the classification process as described in RQ1 for different data granularities. We simulate the different granularities by aggregating the existing 15-minute smart meter data up to the following typical aggregation levels: 15-minute, 30-minute, 60-minute, daily NT/HT (low-tariff during the night / high-tariff during the day), and daily. The NT consumption is measured during different times by different utility companies. For this case we will assume that NT consumption is measured between 23:00 and 07:00.

For each data granularity, we adapted the defined features that are reasonable for the data granularity: For 30-minute and 60-minute data we use features analogously to the 15-minute ones. For the NT/HT and daily aggregation levels we define 14 and 7 different features respectively that describe the consumption during the weekdays and weekends, the relations between different consumptions and the variance. Naturally, a large set of the 15-min smart meter features cannot be calculated for HT/NT or daily measurements (e.g., features on consumption during times of the day, or peaks of the load curve). Additionally, we compare the results to the worst case of only having a single weekly value (1 feature).

The classification is repeated for a single week only. We exclude the weather variables from this analysis, because we want only to show the value of the data at different granularities. The consideration of weather will have more value for the finer granularities, since it is also available on the 15-minute level. For calculations at each granularity levels, the features from coarser granularities are also included. This is done to ensure that important features are not missed on any granularity level. E.g., if there is a good feature that we calculate based on the daily data, but do not for hourly data, then we will get better results for daily granularity than for the hourly data, even though we could get the same or better result by including this feature with hourly data. In this way, we get a large number of features, especially for the 15-minute data, and therefore we perform feature selection with the three best performing methods on the 15-min data (cfs, chi.squared, gain.ratio), and without feature selection.

In Figures 4 and 5, we show the best classification performance that was achieved for different granularities with one of three feature selection methods or without feature selection using Random Forest classifier. Since our goal in this analysis is to compare the best possible performances of different data granularities, we do not create overfitted models that are not necessarily designed to predict household classes for new households, and use the same data for training and test. Therefore, the results cannot be compared with those presented in Fig. 2 and 3, but they give an impression on the impact of data granularity on the classification performance. We can conclude from the results, that ***there is not much difference between using the 15-, 30- and 60-minute data*** for the prediction of the energy efficiency relevant household characteristics. There is a large drop in performance by using the daily (HT/NT and 24-hour) data for some properties. Using the weekly values shows the worst performance, demonstrating the value of finer granularity data.

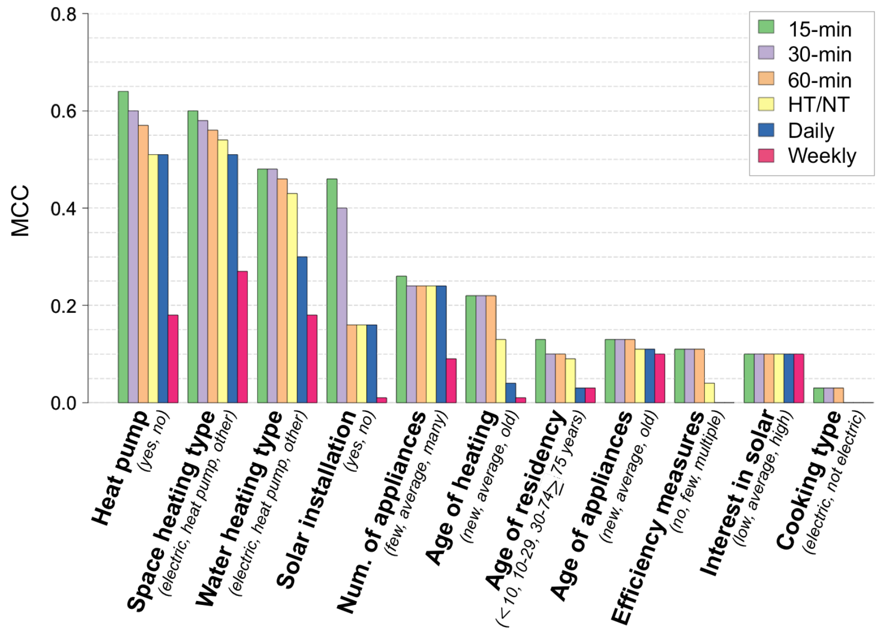

Figure 4: MCC results for different data granularities

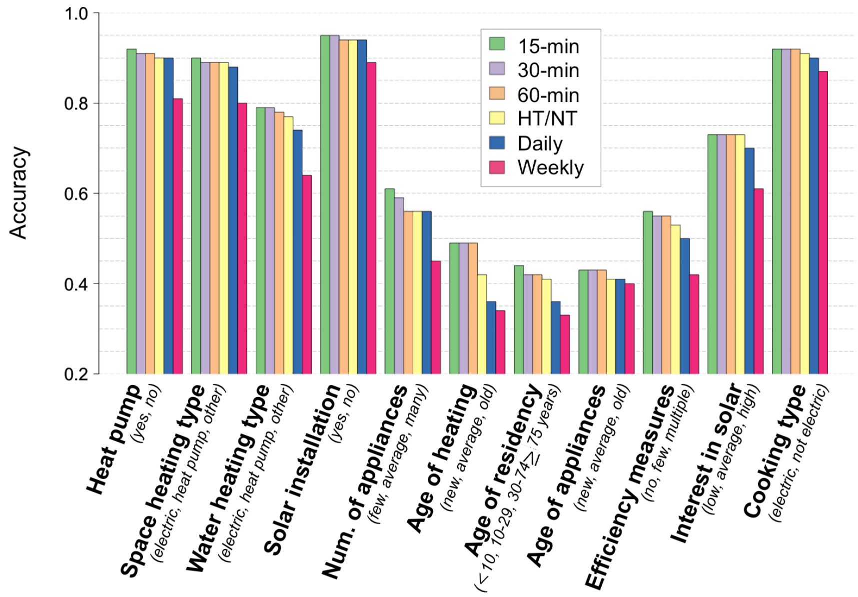

Figure 5: Accuracy results for different data granularities

5 Discussion and Conclusions

The objective of our study was the identification of energy-efficiency and renewable energy related household characteristics based on smart electricity meter data. For the development of our methodology and the training of our machine-learning based green IT artifact, we used smart meter data collected at 15-minute granularity from a Swiss utility company and conducted a customer survey in 2015. In this paper, we have shown that it is possible to identify 9 from 11 proposed energy efficiency relevant household properties (including space heating type and age, water heating type, existence of heat pumps, number and age of appliances, existence and interest in solar installations, and house age) with average accuracy over all properties of 70%. The artifact could not identify the properties "number of completed energy efficiency measures" and "type of cooking facility". Furthermore, we have shown that the smart meter data aggregated at hourly values is sufficient to detect 8 from 11 properties with only a small loss of prediction quality, compared to 15-minute data. Importantly, tools for the recognition of household characteristics must be used under clearly documented privacy preservation conditions and with user consent. Development of law-based guidelines for energy consumption data treatment in the analytics tasks is an important task that should be solved in the next future.

We have identified the following limitations of our work. First, this study is based on the online survey implying the selection bias, which means that we cannot be sure that our sample used for the algorithms training and test is representative and that the results hold for all utility customers. Further, we worked together with a utility company that serves customers from one city and the surrounding area - this introduces a regional bias into the analysis. In addition, we only considered 11 household properties in this work - the results may differ when new properties are included.

In the future work, we plan to conduct field studies to investigate the effects of personalized interventions toward the households selected using the presented algorithms. The interventions can include offers of energy-efficiency products and services, customized consultancy on energy efficiency measures, or normative feedback. We could show the economic, ecologic and social potential of such interventions in a recent study [20] where households with inefficient heating systems were identified based on annual gas consumption data. We expect much better results with the use of smart meter consumption data as shown in this paper. In this setting, we also plan to recognize combined properties (e.g., old house without conducted energy-efficiency measures, or homeowners that are interested in solar installations) from the electricity consumption and weather data. We also plan to reproduce and expand the presented results by cooperating with utility companies serving customers from other geographic locations. A more detailed approach would investigate how the combined changes in granularity and data volume affect the classification performance. Additionally, we used only a simple ensemble learner, that computes the mean value from individual predictions, to aggregate classification results from multiple weeks. Using a more advanced approach that takes the varying performance during different seasons into account could further improve the results.

Being an example for a data science research, our work still allows for empirical validation of the effects from applying the developed artifact in field. Moreover, we demonstrate how to integrate the end-users (utility customers) in the IS research. Ultimately, the proposed artifact is applicable to virtually every smart meter deployment worldwide without changes in the hardware, and thus can ***considerably contribute to the society's energy targets***.

6 Acknowledgements

The research presented herein was financially supported by the Swiss Federal Office of Energy (Grant SI/501202-01) and EUROSTARS (Grant E!9859 - BENgine II).

References

1. European Commission: Commission Recommendation of 9 March 2012 on Preparations for the Roll-out of Smart Metering Systems, http://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32012H0148, (2012).
2. Darby, S.: The Effectiveness of Feedback on Energy Consumption. University of Oxford (2006).
3. Berg Insight AB: Smart Metering in Europe. , Sweden.

4. Degen, K., Efferson, C., Frei, F., Goette, L., Lalive, R.: Smart Metering, Beratung oder Sozialer Vergleich: Was beeinflusst den Elektrizitätsverbrauch? 232 (2013).

5. McKenna, E., Richardson, I., Thomson, M.: Smart Meter Data: Balancing Consumer Privacy Concerns with Legitimate Applications. Energy Policy. 41, 807–814 (2012).

6. Graml, T., Loock, C.-M., Baeriswyl, M., Staake, T.: Improving Residential Energy Consumption at Large Using Persuasive Systems. In: Proceedings of the 19th European Conference on Information Systems (ECIS). AISeL, Helsinki, Finland (2011).

7. Loock, C.-M., Staake, T., Landwehr, J.: Green IS Design and Energy Conservation: an Empirical Investigation of Social Normative Feedback. In: Proceedings of the 32. International Conference on Information Systems (ICIS). , Shanghai, China (2011).

8. Allcott, H., Mullainathan, S.: Behavior and Energy Policy. Science. 327, 1204–1205 (2010).

9. Watson, R.T., Boudreau, M.-C., Chen, A.J.: Information Systems and Environmentally Sustainable Development: Energy Informatics and New Directions for the IS community.(Essay). MIS Q. 34, 23 (2010).

10. Watson, R.T., Howells, J., Boudreau, M.-C.: Energy Informatics: Initial Thoughts on Data and Process Management. In: Brocke, J. vom, Seidel, S., and Recker, J. (eds.) Green Business Process Management. pp. 147–159. Springer Berlin Heidelberg (2012).

11. Birt, B.J., Newsham, G.R., Beausoleil-Morrison, I., Armstrong, M.M., Saldanha, N., Rowlands, I.H.: Disaggregating Categories of Electrical Energy End-Use from Whole-House Hourly Data. Energy Build. 50, 93–102 (2012).

12. Hart, G.W.: Nonintrusive Appliance Load Monitoring. Proc. IEEE. 80, 1870–1891 (1992).

13. Kim, H., Marwah, M., Arlitt, M., Lyon, G., Han, J.: Unsupervised Disaggregation of Low Frequency Power Measurements. In: Proceedings of the 2011 SIAM International Conference on Data Mining. pp. 747–758. Society for Industrial and Applied Mathematics (2011).

14. G. Chicco, Overview and Performance Assessment of the Clustering Methods for Electrical Load Pattern Grouping, Energy 42 (1) (2012) 68–80.

15. Fei, H., Kim, Y., Sahu, S., Naphade, M., Mamidipalli, S.K., Hutchinson, J.: Heat Pump Detection from Coarse Grained Smart Meter Data with Positive and Unlabeled Learning. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1330–1338. ACM, New York, NY, USA (2013).

16. Beckel, C., Sadamori, L., Staake, T., Santini, S.: Revealing Household Characteristics from Smart Meter Data. Energy. 78, 397–410 (2014).

17. Hopf, K., Sodenkamp, M., Kozlovskiy, I., Staake, T.: Feature Extraction and Filtering for Household Classification based on Smart Electricity Meter Data. In: Computer Science-Research and Development. pp. 141–148. Springer Berlin Heidelberg, Zürich (2014).

18. Sodenkamp, M., Hopf, K. and Staake, T.: Using Supervised Machine Learning to Explore Energy Consumption Data in Private Sector housing. In: Handbook of Research on Organizational Transformations through Big Data Analytics, pp.320-333 (2014).

19. Sodenkamp, M., Kozlovskiy, I., Staake, T.: Supervised Classification with Interdependent Variables to Support Targeted Energy Efficiency Measures in the Residential Sector. Decis. Anal. 3, (2016).

20. Kozlovskiy, I., Sodenkamp, M., Hopf, K. and Staake, T.: Energy Informatics for

Environmental, Economic and Societal Sustainability: A Case of the Large-Scale Detection of Households with Old Heating Systems. In: Proceedings of the Twenty-Fourth European Conference on Information Sytsems, AISeL, Istanbul, Turkey (2016).

21. Sodenkamp, M., Kozlovskiy, I., Staake, T.: Gaining IS Business Value through Big Data Analytics: A Case Study of the Energy Sector. In: Proceedings of the Thirty Sixth International Conference on Information Systems (ICIS), Fort Worth, USA, 13-16 December 2015. AISeL, Fort Worth, USA (2015).

22. Hopf. K., Sodenkamp, M., Kozlovskiy, I.: Energy Data Analytics for Improved Residential Service Quality and Energy Efficiency. In: Proceedings of the Twenty-Fourth European Conference on Information Sytsems, AISeL, Istanbul, Turkey (2016).

23. U.S. National Centers for Environmental Information: Climate Data Online, http://www.ncdc.noaa.gov/cdo-web/.

24. Sodenkamp M., Hopf, K., Kozlovskiy, I., Staake, T. Smart-Meter-Datenanalyse für automatisierte Energieberatungen ("Smart Grid Data Analytics") - Schlussbericht für das Bundesamt für Energie Schweiz (Final project report), (2016).

25. Romanski, P., Kotthoff, L.: FSelector: Selecting attributes. (2014).

26. Freund, Y., Schapire, R.E.: A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. J. Comput. Syst. Sci. 55, 119–139 (1997).

27. Zaki, M.J., Meira Jr., W.: Data Mining and Analysis: Fundamental Concepts and Algorithms. Cambridge University Press (2014).

28. Breiman, L.: Random Forests. Mach. Learn. 45, 5–32 (2001).

29. Vapnik, V.N., Vapnik, V.: Statistical Learning Theory. Wiley New York (1998).

30. Cramer, H.: Mathematical Methods of Statistics. Princeton University Press, Princeton (1946).

31. Jurman, G., Riccadonna, S., Furlanello, C.: A Comparison of MCC and CEN Error Measures in Multi-class Prediction. PloS One. 7, (2012).

32. Gorodkin, J.: Comparing two K-Category Assignments by a K-Category Correlation Coefficient. Comput. Biol. Chem. 28, 367–374 (2004).

33. Dietterich, T.G.: Ensemble Methods in Machine Learning. In: International Workshop on Multiple Classifier Systems. pp. 1–15. Springer (2000).