

Submeter based Training of Multi-class Support Vector Machines for Appliance Recognition in Home Electricity Consumption Data

Marco Mittelsdorf¹, Andreas Hüwel², Thole Klingenberg² and Michael Sonnenschein²

¹*Department of Computing Science, University of Oldenburg, Oldenburg, Germany*

²*OFFIS - Institute for Information Technology, Oldenburg, Germany*

Keywords: Appliance Recognition, Smart Metering, Submetering, Energy Monitoring, Multi-class Support Vector Machines.

Abstract: In this paper we employ smart meter and support vector machines (SVM) for the problem of recognizing household appliances' load patterns in measured load time series, which is an important step for various applications in energy consulting, process recognition or health care applications. We present an automated data collection and preprocessing approach that intrinsically avoids many privacy (and security) issues by keeping the whole process local to the household. In the experimental part we investigate multi-class SVMs in the problem domain of automatically recognizing appliances in load profiles of smart meters. For the learning phase, we use low intrusive submeters to automatically and locally generate household specific test data for the supervised training and validation of the SVMs. We analyze classifiers w.r.t. various training sets and feature spaces. Comparing data from household simulator and real household data, we find that excellent recognition rates can be achieved even with low resolution data and rather unsophisticated feature space.

1 INTRODUCTION

The energy transition towards Smart Grid forces energy consumers to adapt to the capabilities of energy producers. This also effects private households. From a more or less passive role that can be characterized by standard load profiles they need to become more aware of the effects of their own energy consumption.

Thus a major goal of all customer information, feedback or consulting systems for electricity consumption is to motivate and deepen the residents understanding of energy consumption, and to possibly trigger investments in energy efficiency or even start changes in behaviour towards a smarter consumption. As suggested by Raabe, Sonnenschein, Beenken, Hüwel and Meinecke (2012) an energy consulting system for private households should give feedback and hints on how to reduce the overall energy consumption on the level of appliance usage. Smart metering is widely discussed for data acquisition in such scenarios. While the global aim is to give insight into power consumption of individual or grouped household appliances, real world scenarios often have the restriction of keeping data acquisition and processing privacy-compliant. This often means that only those aggregated meter data are to leave the

household, that are strictly necessary in respect of invoicing. Any further data, for example needed for the consulting system, must not be transferred somewhere else. The locally gathered total load of the household must be disaggregated into the individual consumption values of each appliance, which is in the fields of Non-Intrusive Appliance Load Monitoring (NIALM). So, in this paper one focus lies on locally labelling the smart meter data, needed for the supervised training phase of our appliance recognition. Another focus is to perform an analysis of different scenarios, how feature space and sampling rate affects event detection and classification.

The rest of this paper is structured as follows. In section 2 we give a short introduction to related approaches in NIALM and in appliance recognition using support vector machines (SVM). In section 3 we present the methodological background that our appliance recognition approach builds on. In section 4 we present our appliance recognition approach and in section 5 we describe our evaluation scenarios and the results of our classifiers. In section 6 we summarize our approach and the results.

2 RELATED WORK

Zeifman, Akers and Roth (2011) give an exhaustive overview over various NIALM approaches that were undertaken between the introduction of the concept by Hart (1992) and today. In this work we focus on supervised learning approaches in order to recognize appliance switching events. For that we need a training phase to train classifiers using labelled data. The quality of such approaches depends strongly on reliable training sets (regarding the labels), the choice of feature space and time resolution of the underlying data.

Following Hart (1992) the time resolution of considered load series must at least be in the range of individual usages of appliances for the purpose of appliance classification. This is mostly a few seconds to subseconds. Some systems even sample at several kilohertz, like the one of Leeb, Shaw and Kirtley (1995), the one of Matthews, Soibelman, Berges and Goldman (2008) or the one of Zeifman and Roth. Their approaches need sophisticated sensory but compared to 1 Hz solutions they enable a new stack of means to the classification problem, which is not focused in this paper.

When using NIALM one major question is how to create the training data and the classifiers. Pihala (1998) gathered his ex ante data of few types of larger appliances over several years in separate field studies, thus once productive the classifiers cannot be further adopted to a specific household. The feedback system of Mattern, Staake and Weiss (2010) allows the consumer himself to manually perform point-wise measurements of a single household appliance, like for instance the power consumption of the computer getting switched on. Lacking an automated approach, the training phase still must be done manually by the consumer himself and this is easily getting tedious and error prone.

When looking beyond the "few large" standard household appliances the system must be able to update the training data and retrain the classifiers. This can either be done by data exchange to the outside of the household or to locally sample and retrain. While Weiss, Staake, Mattern and Fleisch (2012) suggest a system that offensively goes public, real world field scenarios often must be compliant to a more restricting privacy policy, where only those aggregated data may be exchanged to the utility, that are needed for invoicing.

In the context of appliance recognition support vector machines (SVM) can be used to conduct a classification of whether a given switching event was caused by a certain household appliance or not. Based upon training data the SVM constructs a hyperplane,

which separates the items of two classes and allows to classify new observations by simply determining on which side of the hyperplane it lies.

In previous NIALM systems SVMs have mostly been applied to data that was measured using sensors with high sample rates of at least several kHz. To our knowledge Onoda, Murata and Ratsch (2002) were the first to employ SVMs for the task of estimating the state of electric household appliances based on harmonic information. Patel, Robertson, Kientz, Reynolds and Abowd (2007) computed the Fast Fourier Transformation of transient noise signals and used it as a feature. They did also employ SVMs for classification. Lin (2011) stressed that the NIALM problem is in fact a multiple-class decision problem. Very recently Jiang, Luo, and Li (2012) applied multi-class support vector machines (MC-SVMs) to the NIALM problem.

Our approach aims at typical smart meters, which usually have to make use of low cost hardware in order to enable large scale rollouts. This results in rather low sample rates in the range of 15 minutes to one second. So instead of harmonic features we rely on the steady-states and transient variations of the electric load that can be measured with such low frequency sensors.

Kramer et al. (2012) demonstrated that appliance recognition based on such low frequency features can be solved with ensembles of MC-SVMs and K-Nearest Neighbor (KNN) classifiers. They achieved recognition rates of around 95% with the ensemble classifier on a test set consisting of 15 appliances. Furthermore they show that the ensemble classifier outperforms a MC-SVM based on a RBF kernel which yielded recognition rates from 90.5% to 94.3% on the same dataset.

3 DATA ACQUISITION

Our system is designed to stay compliant to a strict privacy policy, which would not allow any exchange of non aggregated power consumption or training data. To gather the needed training data under such determining factors, we adopt an appliance recognition approach similar to the NIALM approach by using low intrusive submeters during the needed training phase. Avoiding error prone manual labelling, they automatically create our training data keeping the whole process local to the household.

Before we introduce our appliance recognition approach we introduce the data our experimental study is based on. We employ two kinds of data sets. First is data that contains measurements of everyday appli-

ances which are turned on or off. We consider this our validation data. We use it to evaluate whether our appliance recognition algorithm performs well on real-world data. Second is data that was generated by our simulation approach. We use this data to implement and test our appliance recognition algorithm.

3.1 Real-world Data

Submetering is one possible approach to create training sets on a per household basis. As the name suggests each appliance is connected to an additional power meter, which enables us to automatically decide whether the appliance is currently switched on or off. We implemented an algorithm which executes the following steps automatically:

1. detect switching events within the data of submeters and generate labels according to the appliance connected to the respective submeter
2. detect switching events within the smart meter data
3. match submeter events and smart meter events by timestamp
4. assign labels of submeter events to matching smart meter events

For details regarding our implementation or the performance of our matching algorithm please refer to (Klingenberg, 2010). The real-world data we use in this paper originated from such a submetering setup. During four weeks we employed the approach for monitoring a small group of six kitchen appliances. The smart meter we used has a resolution of approximately 17 Hz and a basic accuracy of 0.5% of full scale (300 V, 15 A).

3.2 Simulation of Household Appliances

In order to evaluate the results of our appliance recognition approach under various circumstances we employ a simulation approach. Therefore we implemented a household simulator which generates the total load profile of a model household as well as a load profile of every single simulated appliance. This is the same data which we acquire by employing the abovementioned submetering approach. In order to further improve the comparability of both approaches the simulator does also generate data at a resolution of 17 Hz.

The easiest way to create a simulation model of an appliance, that reproduces steady state transitions as well as the transient behaviour during switching events realistically, is to record load patterns during appliance operation. At simulation time the appliance

model steps through one, e.g. randomly chosen, pattern value by value.

As was stated by Hart (1992), Pihala (1998), Baranski (2006) and others the electric load is highly dependent on the fluctuating voltage signal. Electric admittances though are not influenced by the voltage signal in theory. Admittance values are therefore a better representation for our appliance patterns than active and reactive power values. We do therefore use recorded voltage signals $U_{a,t}$ and current signals $I_{a,t}$ of an appliance a to compute the admittance

$$Y_{a,t} = \frac{I_{a,t}}{U_{a,t}} \quad (1)$$

for each time step t of the signal and use it to compute the values of conductance $G_{a,t}$ and susceptance $B_{a,t}$ by applying the following relation :

$$\underline{Y}_i = G_{a,t} + j \cdot B_{a,t} \quad (2)$$

$$= Y_{a,t} \cdot (\cos \varphi - j \cdot \sin \varphi) \quad (3)$$

which yields the calculation rules

$$G_{a,t} = Y_{a,t} \cdot \cos \varphi \quad (4)$$

$$B_{a,t} = -Y_{a,t} \cdot \sin \varphi \quad (5)$$

where $\varphi = \varphi_u - \varphi_i$ denotes the phase angle between current and voltage signal. This way we prepared a pattern comprised of time series G_a and B_a for each appliance we want to include into our simulator.

For simulation purposes we employ a simplified electric single-phased model of households. We basically assume that all appliances are plugged together in a parallel connection. This enables us to compute the total admittance of the household:

$$\underline{Y}_{tot} = \sum_{i=1}^N \underline{Y}_i \quad (6)$$

At simulation time a predefined operation schedule specifies when appliances are switched on or off. The simulator does also generate a fluctuating voltage signal U_0 which we use to calculate the total active power P and reactive power Q of the household and of each appliance according to the following formulas:

$$\underline{I} = \underline{U} \cdot \underline{Y} \quad (7)$$

$$\underline{S} = \underline{U} \cdot \underline{I}^* \quad (8)$$

$$P = \text{Re}(\underline{S}) \quad (9)$$

$$Q = \text{Im}(\underline{S}) \quad (10)$$

4 OUR APPLIANCE RECOGNITION APPROACH

Our approach is divided into the following three main steps – here presented for the virtual household, for the real household we use smart meter and submeter.

1. Training of classifiers: We use the household simulator to generate labelled training data for a given set of appliances and train a MC-SVM.
2. Scenario classification: We use the household simulator to generate a new total load profile, consisting of the same set of appliances and use the MC-SVM to classify the detectable switching events.
3. Scenario evaluation: We compare the results of the classification process to the load profiles of the appliances and generate a summarizing statistic about matches and mismatches.

One main component of steps 1 and 2 is our event detection and feature extraction algorithm. Therefore we will start with a description of this algorithm before we describe the three main steps in the following sections.

4.1 Event Detection & Feature Extraction

Our event detection algorithm is based on the active power signal and consists of the following six steps and an optional seventh step. The first step is a pre-processing operation according to the suggestion of (Hart, 1992) to only use normalized power values for appliance recognition. We therefore eliminate the influence of the fluctuating power signal by applying

$$P_{\text{normalized}} = U_{\text{ref}}^2 \cdot G \quad (11)$$

Where U_{ref} denominates the nominal value of the supply voltage and G is the electric conductance.

In the second step we compute the series of differences of this normalized power time series:

$$\Delta P_t = P_t - P_{t-1} \quad (12)$$

According to (Baranski, 2006) this representation is well suited for the detection of switching events because they do appear as peaks whereas steady-state segments of the load profile appear as values around zero.

In the third step we apply a *global* filter to the series of differences in order to remove most of the noise events from the signal. This global threshold is to be carefully chosen because if it is too high it suppresses events of some appliances. If it is chosen too low a significantly higher amount of noise events has to be handled during recognition phase.

In the fourth step continuous sections of positive or negative values within the series of differences are combined. This enables us to combine staircase-shaped event sequences to one single event. We found that most staircase-shaped events have a duration of

Table 1: Overview of the features we use for classification.

Feature	Description
ΔP	Change in active power
ΔQ	Change in reactive power
ΔZ	Change in impedance
ΔR	Change in resistance
ΔX	Change in reactance
ΔY	Change in admittance
ΔG	Change in conductance
ΔB	Change in susceptance
P_{Surge}	Maximum active power of the surge
t_{Surge}	Duration of the surge

less than one second. The maximal duration of a section is therefore limited to one second. It follows that events of different appliances need to be at least one second away from each other.

In the fifth step all events of a section are added up. We use the sum of the switching powers to represent the total power of the combined switching event.

The sixth step is only necessary for extracting accurate training data and is skipped during the classification of a scenario load profile. In this step an appliance specific filter is applied to separate noise events from events that represent a relevant state change. Noise events may occur during appliance operation. They often have a lower power draw than on and off events but they exceed the global filter threshold. Switching on power surges are treated as two events. One indicates the rising edge and the other one the falling edge of the surge. With an optional seventh step a surge event can be detected and combined to a single event.

Other electric features such as reactive power, admittance or resistance are extracted based upon the sections which were identified in step four and maybe further combined in step seven. The change of these features is computed by subtracting the value at the begin of a section from the value at the end of the section. Table 1 gives an overview of the features we use for classification.

4.2 Training of Classifiers

Supervised machine learning algorithms like SVM need labelled training data. In order to generate labelled data of a given appliance we use submeters or our simulator to get load profiles containing switching events of only this appliance. Next we apply our event detection algorithm. The events of each appliance are divided into on and off events and labelled respectively. Noise events which do not represent state changes of appliances are also divided into events with ascending power change and events with

descending power change. In total $2n + 2$ (2 noise classes and $2n$ classes for all appliances) distinguishable labels are generated where n is the number of appliances considered.

4.3 Scenario Classification

First step of the classification is to run the event detection and feature extraction algorithm on the load profile to gain appropriate input data for the MC-SVM. In the second step every event is classified.

4.4 Scenario Evaluation

During this phase we evaluate how well the classifier performed during the aforementioned classification phase. To find the true class of an event we use the appliance load profiles which are also provided by the household simulator. The event detection and feature extraction algorithm is used to determine the switching events within these load profiles. Next we compare these events with those that were classified earlier. If timestamp and label are equal then the classification was correct. If the timestamp matches but the label does not then a misclassification occurred.

5 EXPERIMENTAL STUDIES

In the first part of this section we introduce three scenarios to test and to evaluate our system on. In the second part we present the results of each scenario. A full description of the scenarios and results can be found in (Mittelsdorf, 2012).

5.1 Scenarios

Scenario 1 - Feature Space. Here we analyze the influence of the feature selection on the appliance recognition rates. Out of the available 27 appliances of the household simulator we chose a subset of 18 appliances for the base scenario (1a). Two additional appliances are considered in two variants of the base scenario: a fridge (1b) and a washing machine (1c). The seven remaining appliances are either not representative because we could not capture enough data or the event detection does not work reliably for them.

The power draw of a stereo system for instance heavily depends on the current volume level and the song just played. (Baranski, 2006) counts the stereo system to continuously varying electric consumers which are much more demanding than simple steady state on-off-appliances with constant power draws. Another appliance which we do not consider is for

example the 20 W desk lamp. Since a global 40 W filter led to the best results for the majority of appliances, switching events of smaller consumers are filtered out.

To create scenario 1a we simulated a load profile with a duration of 24 hours. All appliances were scheduled to switch on and off several times and overlap in their operating time. Altogether 1230 events were detected. Figure 1 shows that some of the appliances have a similar power draw. We presume that by using only the active power as feature misclassifications will occur. Our assumption is therefore that classification rates will improve by adding more features.

Scenario 2 - Sampling Rate. In order to investigate the influence of the sampling rate on appliance recognition we use the same scenario with a 17 Hz and a 1 Hz sampling rate. The latter one is created by down-sampling the load profile of the 17 Hz scenario. The appliance set of this scenario consists of an incandescent lamp, a water kettle, a vacuum cleaner and a TV. The lamp has a low and the water kettle a high power draw. The TV and the vacuum cleaner show a power surge in their load profiles whereas the other two appliances do not. The set therefore covers appliances with different characteristics.

Scenario 3 - Real Data. In order to investigate the performance of our system applied to real world data we compare the classification results achieved on the real world data mentioned in section 3.1 to the results of simulated load profiles. The real world data we use was recorded by (Klingenberg, 2010) and consists five appliances: dishwasher, fridge, coffee maker, boiler and water kettle.

5.2 Results

Out of the confusion matrix various classification parameters can be calculated. We use the hit rate, false alarm rate, precision and ROC-Score. We extended the classical confusion matrix for two class problems to our multi-class classification approach. Figure 2 shows this for the switch-off class of the fridge. While the true positive rate TP is still a single value false negative FN, false positive FP and true negative rates TN now represent the summarized values of the accordingly row (FN), column (FP) or main diagonal (TN).

Scenario 1 - Feature Space. The results of scenario 1a are presented in Figure 3. It shows the four classification parameters for different features. The upper di-




Figure 1: Distribution of switching power from the appliances of scenario 1. Areas of similar switching power of two or more appliances are highlighted red.

agram shows mean values whereas the lower diagram shows the worst achieved values for a single class. The best results were achieved by using the features (ΔP) and $(\Delta P, \Delta Z)$. The mean hit rate for (ΔP) was 96.25 % and for $(\Delta P, \Delta Z)$ was 96.2 %. The worst hit rate for (ΔP) was 70 % as the lower diagram shows, which means that at least for one class only 70 % of the switching events were classified correctly. Most of the feature combinations used in scenario 1 have satisfying mean results as the upper diagram in Figure 3 shows, but for some combinations worse results were achieved for a single class. The minimal hit rate with the features $(\Delta P, \Delta Q, \Delta Y)$ was 0 % which means switching events of one class haven't been classified correctly at all. Therefore this combination is not suitable for our appliance recognition system. Other unsuitable features are (ΔQ) , (ΔB) and (ΔX) which have a mean hit rate of less than 25 %.

In scenario 1b we added a fridge and in scenario 1c we added a washing machine as additional appliances to the base scenario. The mean hit rate for scenario 1b was 93 % and for scenario 1c only 85.94 %. The worse results of scenario 1c can be explained by the high number of switching events that occur during washing machine operation and their different switching powers. The majority of the washing machine switching events have a switching power in the range of 40 W to 400 W and -40 W to -400 W. Eight out of

the 18 appliances show also switching powers in this range, so it is more likely that events of those appliances are classified as washing machine and vice versa washing machine events are classified as events of those appliances.

Scenario 2 - Sampling Rate. Let us look at the accuracy of the event detection at both sampling rates. In case of the 17 Hz sampling rate the event detection works precisely and detects all events of the four appliances we use in this scenario. Applied to the 1 Hz data, the event detection misses 11 out of 151 events.

a1	b1	c1	d1	e1	f1	g1	h1	i1	j1	k1	l1	
5				1								<= classified as
	4											a1 = Dishwasher_Real_off
		91										b1 = Dishwasher_Real_on
			TP									c1 = Coffee_machine_Real_off
												d1 = Coffee_machine_Real_on
												e1 = Fridge_Real_off
FN		80		401					32			f1 = Fridge_Real_on
				286					12			g1 = Boiler_Real_off
					15							h1 = Boiler_Real_on
						1			14			i1 = Electric_kettle_Real_off
										1		j1 = Electric_kettle_Real_on
												k1 = Noise_off
												l1 = Noise_on
												FP
												TN

Figure 2: Confusion matrix of multi-class classification.

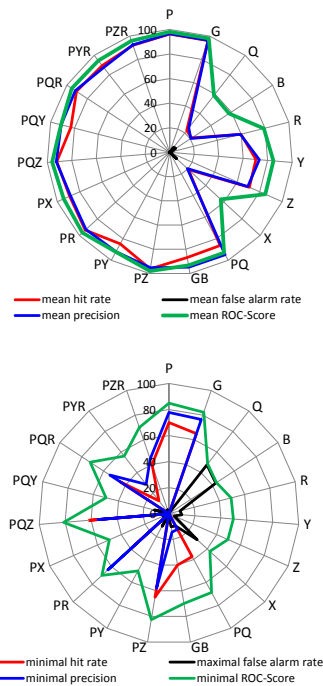


Figure 3: Classification results for different features of scenario 1a in percentage. Mean results of all classes on the top and worst results of a single class on the bottom.

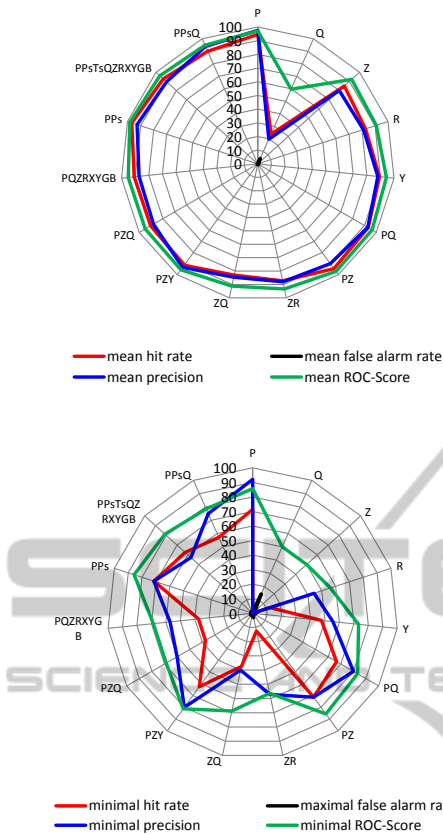


Figure 4: Classification results for different features of scenario 3. Mean results of all classes on the top and worst results of a single class on the bottom.

This leads to an accuracy of 92.72%. If the noise events are also considered then only 87% of the 17 Hz events were detected in the 1 Hz data.

The quality of the classification results also depends heavily on the sampling rate. On the 17 Hz data we could achieve an average hit rate of 90.62%. The hit rate achieved on the 1 Hz data was significantly worse and amounted to 74.64%.

The lower accuracy of both event detection and classification combines to dramatically lower accuracy of the whole appliance recognition system.

Scenario 3 - Real World Data. The best result was achieved with the features $(\Delta P, P_{Surge})$ with a mean hit rate of 96.8%. Also the features (ΔP) and $(\Delta P, \Delta Z)$ led to good classification results as Figure 4 shows. The evaluation of the real world scenario shows similar results as for the simulated data. From this we conclude that our simulator is suitable to generate data close enough to reality.

6 DISCUSSION

In this work we demonstrated how an appliance recognition system can be set up by the use of submetering data. We presented the concept of our household simulator which allows us to design and generate load profiles of a model household and of single appliances for testing purposes. We also presented our appliance recognition algorithm which was tested on such simulated data and on a real world dataset.

We investigated which choice of features leads to the best results in appliance recognition using MC-SVMs. Therefore we performed a feature study on simulated and on a real world scenarios. In both cases most scenarios have shown satisfying recognition results, if the change of active power ΔP is among the considered features. Also the feature combinations (ΔP) or $(\Delta P, \Delta Z)$ which perform best on simulated data did perform very well on the real world data. These results indicate that our simulator is a reasonable model of real households. A more systematic approach to validate our household simulator might be to recreate a real world scenario using simulation models of the exact same appliances that were used in reality. The two resulting load profiles could be compared by calculating the residual sum of least squares, which should ideally have the value zero.

Most appliance recognition systems rely on data with a resolution of several kHz or on data with a resolution of 1 Hz. We investigated how the recognition rates behave if slightly higher resolutions than 1 Hz, such as 17 Hz, are used. Based on the findings from our scenario appliance recognition systems at 17 Hz do significantly outperform systems at 1 Hz. The reason is that the 1 Hz system drops more events during event detection and in addition it recognizes only about 75% correctly during event classification.

On the real world data we were able to achieve a recognition rate of 96.8% with the feature combination $(\Delta P, P_{Surge})$. (Kramer et al., 2012) used the same sensor as we did and achieved recognition rates of up to 95% with ensemble classifiers created from SVM and KNN classifiers. They found recognition rates ranging from 90.5% to 94.3% using a MC-SVM based on an RBF kernel. On the first glance our algorithm performs better, but it should be mentioned that our real world data consists of five individual appliances whereas the data investigated by (Kramer et al., 2012) consists of 15 individual appliances. So the complexity of their classification problem is higher. On the other hand the data investigated by (Kramer et al., 2012) consists of hand-picked, manually labelled patterns and is balanced whereas our dataset was automatically detected within data

from everyday life and automatically annotated using the additional data of submeters. Since we aim for a fully automated approach we additionally have to automate the decision whether an event originated from the state change of an appliance or from noise in the power signal. We incorporated this decision into the classification problem by adding additional classes for noise events.

We addressed privacy indirectly by designing our system in a way that allows us to train classifiers on a per household basis. This means that all personal data is processed in-house and never uploaded or processed anywhere else.

One drawback of our system is, that all classifiers have to be retrained if a new appliance is added to the household. In such a case the data of most appliances can be reused, but during a short setup phase patterns of the new appliance must be gathered. So at least one submeter should permanently be available in the household whereas most other submeters can be removed after the initial setup.

ACKNOWLEDGEMENTS

This work has been supported by funds of the Federal Ministry of Economy and Technology in the E-Energy project *eTelligenz*, project number 01MR08007A.

REFERENCES

- Baranski, M. (2006). *Energie-Monitoring im privaten Haushalt*. PhD thesis, University of Paderborn.
- Hart, G. (1992). Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891.
- Jiang, L., Luo, S., and Li, J. (2012). An Approach of Household Power Appliance Monitoring Based on Machine Learning. *2012 Fifth International Conference on Intelligent Computation Technology and Automation*, pages 577–580.
- Klingenberg, T. (2010). *Smart Submetering - Effizienter Einsatz von Submetern zur Aktivitätsbestimmung in Privathaushalten mit Hilfe adaptiver Lernverfahren*. Master's thesis, University of Oldenburg.
- Kramer, O., Wilken, O., Beenken, P., Hein, A., Hüwel, A., Klingenberg, T., Meinecke, C., Raabe, T., and Sonnenschein, M. (2012). On ensemble classifiers for nonintrusive appliance load monitoring. *7th International Conference on Hybrid Artificial Intelligence Systems (HAIS)*, pages 322–331.
- Leeb, S. B., Shaw, S. R., and Kirtley, J. L. (1995). Transient Event Detection in Spectral Envelope Estimates. *IEEE Transactions on Power Delivery*, 10(3):1200–1210.
- Lin, Y.-h. and Tsai, M.-S. (2011). Applications of hierarchical support vector machines for identifying load operation in nonintrusive load monitoring systems. *2011 9th World Congress on Intelligent Control and Automation*, pages 688–693.
- Mattern, F., Staake, T., and Weiss, M. (2010). ICT for green: how computers can help us to conserve energy. *International Conference on Energy-Efficient Computing and Networking*.
- Matthews, H. S., Soibelman, L., Berges, M., and Goldman, E. (2008). Automatically disaggregating the total electrical load in residential buildings: a profile of the required solution. In *International Workshop on Intelligent Computing in Engineering*, page 381389.
- Mittelsdorf, M. (2012). *Evaluierung des Einflusses von Merkmalsauswahl und Abtastrate auf die Geräteeerkennung mit Support-Vektor-Maschinen*. Master's thesis, University of Oldenburg.
- Onoda, T., Murata, H., and Ratsch, G. (2002). Experimental analysis of support vector machines with different kernels based on non-intrusive monitoring data. *Neural Networks, 2002.*, pages 2186–2191.
- Patel, S., Robertson, T., Kientz, J., Reynolds, M., and Abowd, G. (2007). At the flick of a switch: Detecting and classifying unique electrical events on the residential power line (nominated for the best paper award). In Krumm, J., Abowd, G., Seneviratne, A., and Strang, T., editors, *UbiComp 2007: Ubiquitous Computing*, volume 4717 of *Lecture Notes in Computer Science*, pages 271–288. Springer Berlin Heidelberg.
- Pihala, H. (1998). *Non-intrusive appliance load monitoring system based on a modern kWh-meter*. Technical Research Centre of Finland.
- Raabe, T., Sonnenschein, M., Beenken, P., Hüwel, A., and Meinecke, C. (2012). Energieberatung in Haushalten auf basis des smartmetering. *Ökologisches Wirtschaften*, 1:46–50.
- Weiss, M., Staake, T., Mattern, F., and Fleisch, E. (2012). Powerpedia: changing energy usage with the help of a community-based smartphone application. *Personal and Ubiquitous Computing*, 16(6):655–664.
- Zeifman, M., Akers, C., and Roth, K. (2011). Nonintrusive appliance load monitoring (nialm) for energy control in residential buildings. *International Conference on Energy Efficiency in Domestic Appliances and Lighting (EEDAL), Copenhagen*.
- Zeifman, M. and Roth, K. (2011). Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics*, 57(1):76–84.