



Insights into Unsupervised Holiday Detection from Low-Resolution Smart Metering Data

Günther Eibl^(✉), Sebastian Burkhart, and Dominik Engel

Center for Secure Energy Informatics, Salzburg University of Applied Sciences,
Urstein Süd 1, 5412 Puch/Hallein, Austria

{guenther.eibl,sebastian.burkhart,dominik.engel}@en-trust.at
<http://www.en-trust.at>

Abstract. Recently, first methods for holiday detection from unsupervised low-resolution smart metering data have been presented. However, due to the unsupervised nature of the problem, previous work only applied the algorithms on a few typical cases and lacks a systematic validation. This paper systematically validates the existing algorithm by visual inspection and shows that numerous cases exist, where implicit assumptions are not met and the methods fail. Moreover, it proposes a new, very simple rule-based method which is in principle able to overcome these problems. This method should be seen as a first step towards improvement, since it is not automated and needs a moderate amount of human intervention for each household.

Keywords: Privacy · Smart grids · Smart metering

1 Introduction

Privacy concerns have been raised by the planned large-scale rollout of Smart Meters [10]. Several methods exist that analyze energy consumption profiles. Especially NonIntrusive Load Monitoring (NILM) is an area of active research [4,7,11] where the goal is to divide (disaggregate) the (aggregate) signal that results from summing up consumption of different devices into the individual summands. While it is clear that such techniques have high potential to improve or generate new services, one could learn private information like habits from information about device usage. Typically, these approaches use a time resolution in the range of a second which is much lower than the time resolutions that occur in the Smart Grid. For example in Austria, the planned rollout foresees Smart Meters that only have a time resolution of 15 min.

Another branch of methods aim at detection of occupancy, i.e., at determining for each point in time, if a household is occupied by at least one person or not [1,2,5,8,9]. These methods are typically tested using laboriously collected ground truth for a low number of places. Also there, typically time resolutions of seconds or a minute are smaller than the time resolution in Smart Metering.

In a recent study [3] holiday detection from low resolution has been investigated for the first time. Inspired by ideas from occupancy detection several methods have been developed that aim at detecting holidays. Note that there and in this paper a holiday is defined as a day without electricity consumption that is triggered by a human. The methods have been described and investigated for several typical households. However, since the paper was intended as a first, explorative step towards holiday detection, no systematic validation of the methods has been performed.

This paper aims at filling this gap. The performed validation discovered several weaknesses of the existing methods and proposes a first step towards a better method.

2 Related Work

2.1 NIOM

NIOM is a rule-based algorithm that classifies households as occupied or unoccupied for time intervals containing T measurement points. For each time interval the following features are calculated: the average value, the standard deviation and the range (maximum-minimum). For each feature, thresholds need to be found for classifying the time interval as unoccupied, if the feature is below the threshold. Otherwise the time interval is classified as occupied. All the time intervals where occupancy is determined that way occur during daytime $\mathcal{D} = [6 \text{ h}, 23 \text{ h}]$ and not during nighttime.

The thresholds themselves are determined as the maximum values during the previous night between 1 h and 4 h ($\mathcal{N} := [1 \text{ h}, 4 \text{ h}]$). Gaps between time intervals detected as occupied are filled using a plausible heuristic. The authors themselves already pointed out that their results are sensitive to the choice of the threshold.

2.2 MaxOnly

MaxOnly has been developed in [3] as a simple extension of NIOM for holiday detection. As done by NIOM, the maximum value of the measurements of household i , but now only one value per 15 min, is determined for the nighttime \mathcal{N} of day d . Then the maximum of the day values $x_{i,d}^t$ is compared with this value. In order to not underestimate the number of holidays a tolerance δ is added to the night maximum yielding the *MaxOnly* holiday detection rule

$$h_{i,d} = \begin{cases} 1 & \text{if } \max_{t \in \mathcal{D}} x_{i,d}^t < \max_{t \in \mathcal{N}} x_{i,d}^t + \delta \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

For practical tests, δ has been set to 0.1 kW.

2.3 LogReg

While MaxOnly is a simple extension of NIOM, LogReg has been developed in [3] especially for holiday detection. The idea of LogReg is to reformulate the problem as follows: For a holiday, both, the night and the day values of a day, should resemble night values. This reformulation enables one to train and use a classifier f_i that distinguishes night values from day values of household i . Thus, in a first step a classifier f_i is learned for each household using logistic regression which is a common classification method. Logistic regression has been used because it not only outputs the class but also the estimated probability to be in a class.

Given the learned classifier, a new day to be classified is then treated as follows: if the classifier classifies both, its day values and its night values, as likely being night values using again a threshold τ_i then the day is classified as a holiday. The threshold τ_i is heuristically determined from the 25% quantile of the estimated probability of the night measurements to be a night measurement.

So far it is still open which features should be used as the input to the classifier, i.e., how to describe the day and night values, respectively. Two variants have been studied in [3] and will also be studied here. The first set of features summarizes the day and night values simply by their maximum, the corresponding overall method is called *LogRegMax*. In order to possibly exploit more details of the distribution the second set of features summarizes the day and night values by quantiles. More precisely the 25%, the 50%, the 75% and the 99.9% quantile are used. The corresponding overall method is called *LogRegQuant*.

3 Validation

3.1 Initial Comparison of Results

It has already been stated in [3] that the differences in the predicted number of holidays can be large. Now, it is investigated where this differences come from and which of the methods yields the better prediction. The proceeding for this task is as follows: First the households where the largest differences in the




Fig. 1. Comparison of the results of the methods proposed in [3].

predictions occur are detected, then validation plots were created and visually compared. Essentially the result of this section is a complete failure of all the existing methods for a variety of households.

By looking at the distribution of the number of holidays, it can be shown that NIOM and LogRegQuant yield to lower number of holidays than MaxOnly and LogRegMax (Fig. 1).

It is clear that NIOM yields fewer holidays than MaxOnly because its threshold is lower by the chosen tolerance (chosen as 0.1 kW). We explain the underestimation by LogRegQuant by the conjecture that the maximum is the best feature for the description of human activities. More details will be shown in Sect. 5.1.

Preliminary examinations of validation plots provided first evidence that MaxOnly and LogRegMax are favoured over NIOM and LogRegQuant. Consequently the subset of households that are investigated are based on the differences between these two methods which is shown in Fig. 2. In order to easily recognize households where differences occur the points of the scatterplot are labeled by the household ID.

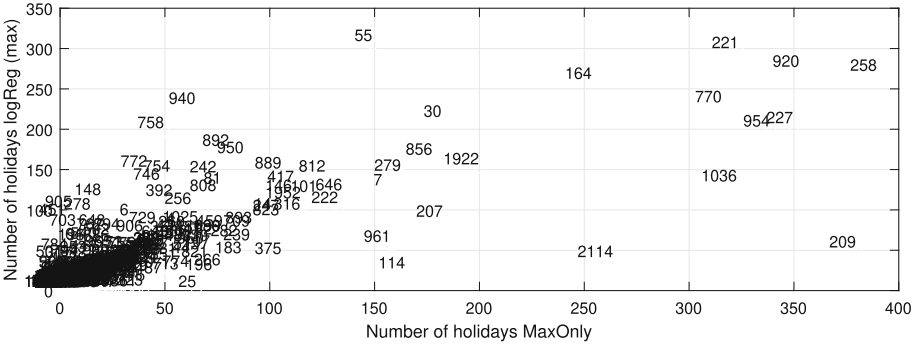


Fig. 2. Comparison between the results of MaxOnly and LogRegMax.

3.2 Validation Plots

Since no information about the true holidays exists human inspection must serve as the gold standard here. Therefore, validation heavily relies on visual analysis and common sense. In order to produce a plot that quickly enables a human to estimate the plausibility of the result for each household and method the days are ranked according to the quantity that represents the belief in being a holiday of the respected method. After reordering the days according to this belief the values of the household are again plotted as a heatmap. This heatmap is then visually analyzed for plausibility. As described below, although this validation method is simple, it enables a human to find implausible results of the investigated methods.

Note that the ranking depends on the prediction method: for NIOM and MaxOnly holiday score s_d^{MaxOnly} is calculated as the according to the difference

between the maximum value during nighttime and the maximum value during the daytime. Note that while the score is the same for NIOM and MaxOnly, the threshold is set higher for MaxOnly.

For the logReg methods the holiday score $s_{i,d}^{\text{logReg}}$ for day d of household i is the predicted probability for day d to be a holiday

$$s_{i,d}^{\text{logReg}} = p_d^i = y_d^{i,\text{day}} \cdot y_d^{i,\text{night}}$$

While the formula for the score is the same for the two variants, the difference lies in the inputs of the trained classifiers f^i that tries to distinguish between the day and the night values.

In order to better visualize the differences between the night and the day values for the heatmap values that exceed 1 kW were set to 1 before visualization. More details about this choice is given in Sect. 5.1. An additional vertical separation line with numbers of 1 kW is added in order to separate days which were detected holidays (left) from days where no holiday was detected (right).

The reordering has an additional beneficial effect. While longer periods of holidays such as in Fig. 3 can easily be detected in the original heatmap, single holidays can easily get overlooked. Due to the reordering, single holidays are put to the left of the separation line and can thus be compared with other detected holidays.

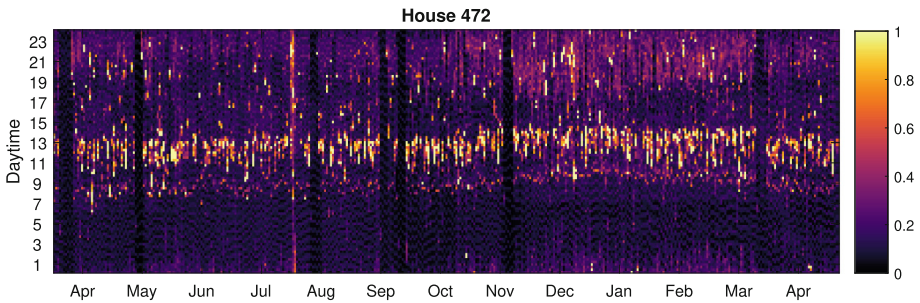


Fig. 3. Household 472: holiday periods can be easily identified as vertical black stripes.

Based on the visual inspection of the validation plots the following categories were found to be suitable for describing the results: “plausible”, “too few holidays detected”, “too many holidays detected” and “bad ranking”. Since the number of holidays is not known the best category to achieve is plausible. This means that the sorting of days in the heatmap and the distinction between holidays and non-holidays look reasonable. “Bad ranking” means that the ordering of days in the heatmap does differs from the expected ordering. An example for bad ranking is shown in the upper panel of Fig. 4. There, many days with high consumption especially during night are rated as most likely being a holiday. This is far from being plausible since during a holiday low consumption is expected throughout

the day. An example for too many found holidays is shown in the lower panel of Fig. 4. While the ranking is plausible there, only the days with consistently very low energy consumption that constitute the black part at the leftmost part of the figure should be labeled as holidays.

Due to the high number of households, not all validation plots were looked at. Instead, in spirit of estimating the performance from properly selected samples four different regimes of households were investigated: two kinds of households, where the two methods agree and two kinds, where the two methods disagree. For the latter two cases households with an extremely high number of predicted holidays by at least one method were investigated. More precisely it is looked at households for which

- both methods estimate around 30 holidays: 7 selected households
- both methods estimate around 0 holidays: 6 selected households
- MaxOnly predicts more than 150 holidays: 18 selected households
- logRegMax estimates more than 150 holidays but MaxOnly estimates less than 100 holidays: 7 selected households

While it is not possible to even show all these plots, typical households are shown in order to demonstrate the key issues.




Fig. 4. Household 892: both MaxOnly (top) and LogRegMax (bottom) overestimate the number of holidays. In the panel above, the false positives of MaxOnly that occur from high night values can be seen most clearly.

3.3 Households with About 30 Detected Holidays

First, cases that may be considered normal for Austria (with 25 free days per year) are analyzed now. A subsample consisting of households 1, 112, 194, 334, 551, 552, 472 was investigated. Except household 551 all of them had at least one longer period with a holiday. Both algorithms showed plausible results with the

one exception when MaxOnly wrongly assigns a day a holiday when larger values occur during night time. As a typical example the consumption of household 472 is shown in Fig. 3.

The validation plots for the same household 472 are shown in Fig. 5. As already described in [3], both methods yield plausible results for such cases: in this case 25 holidays are detected by MaxOnly, 21 by LogRegMax.

3.4 Households with About 0 Detected Holidays

Considering households with near zero assigned holidays it is studied, if the methods also tend to underestimate the number of holidays. Visual inspection of households 73, 422, 452, 460 and 968 confirm that no longer holiday exist and only a small number of single holidays may exist. Thus also for this regime the algorithms yield plausible results.

Out of these households, household 769 is the only one that has a short holiday period (Fig. 6) in January. For this household logRegMax shows a plausible ranking but too few holidays: days with low consumption throughout the day are right to the separation line (Fig. 7). MaxOnly shows more holidays.

This household is also an interesting case where a pool exists. This pool causes high values of a rectangular shape both during day and night. In fact, one of MaxOnly's detected holidays has a considerable part of the day affected by the pool pump. Because the remaining values of this day are small, this may indeed be a day where no one is at home and only the automatic pool pump is running as the only device with high energy consumption. This household demonstrates that detection methods must find a way to deal with automatic appliances like swimming pools.

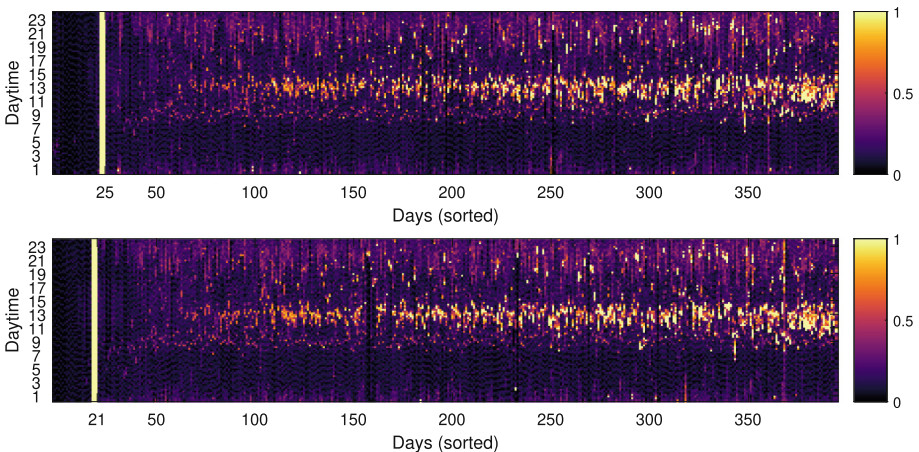


Fig. 5. Household 472: both algorithms show plausible results: MaxOnly (top): 25, LogRegMax (bottom): 21 holidays.

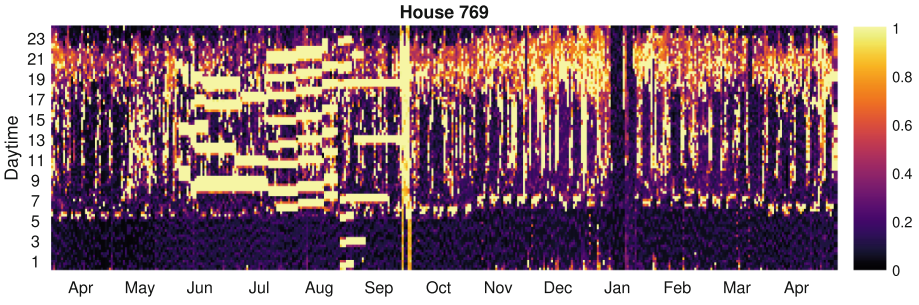


Fig. 6. Household 769 has one holiday in January and rectangular patches of higher values that are caused by a pump of a swimming pool.

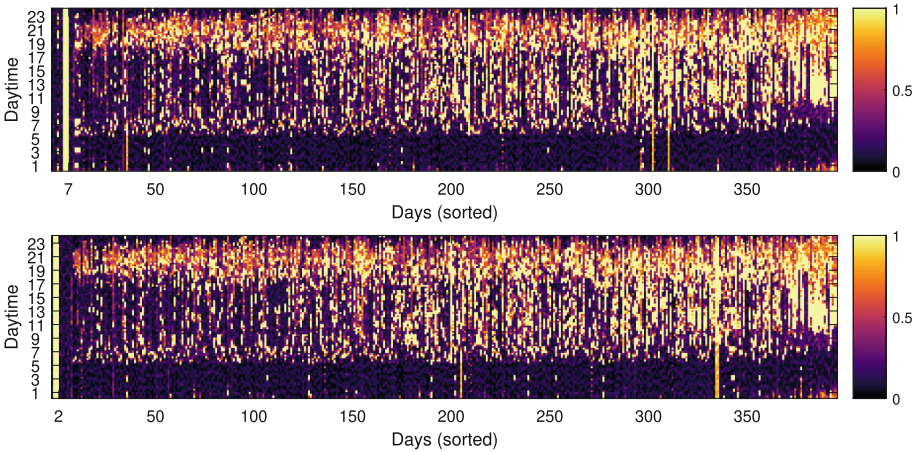


Fig. 7. Household 759: MaxOnly shows a plausible result with one questionable holiday (top), LogRegMax underestimates the number of holidays (bottom).

Table 1. Validation of MaxOnly for households where MaxOnly estimates ≥ 150 holidays.

Validation result	Household
Plausible estimate	114, 207, 227, 258
Bad ranking	7, 30, 164, 279, 770, 856
Estimated too many	221
Estimated too few	209, 920, 954, 961, 1036, 1922, 2114

3.5 Households with More Than 150 Holidays for MaxOnly

Now households where MaxOnly estimated more than 150 holidays are visually validated. The IDs of these households can also be identified in the scatter plot in Fig. 2. The validation result for MaxOnly is shown in Table 1.




Fig. 8. Household 770: the regular, “oscillating” behaviour is supposed to be caused by the refrigerator. This background signal is not there the whole year.

Since the estimated number is so high the authors expected an overestimation of MaxOnly. However, this turned out to be not the case. The validation plots even more showed that more often both methods still underestimated the number of holidays. In the best case results looked plausible, however this did not happen often.

MaxOnly mainly also suffered from an inadequate ranking. Bad ranking of MaxOnly occurs for days with low values during daytime and at least on high consumption value during nighttime. Since MaxOnly compares the maximum during daytime with the maximum during nighttime, the difference is negative which is considered as a holiday by MaxOnly. However, such a day is likely not a holiday resulting in erroneously detected holidays where the consumption during day was low but consumption during night was high. An example household, where bad ranking clearly occurs, is shown in Fig. 12.

The validation result for LogRegMax shown in Table 2 demonstrates that for these households LogRegMax in most cases underestimates the number of holidays.

Table 2. Validation of LogRegMax for households where MaxOnly estimates ≥ 150 holidays.

Validation result	Household
Plausible estimate	7, 164
Bad ranking	221
Estimated too many	30
Estimated too few	114, 207, 209, 227, 258, 279, 770, 856, 920, 954, 961, 1036, 1922, 2114

As an example, household 770 is shown in Fig. 8. The figure clearly demonstrates that the background signal can vary over the time of the year. This property makes the determination of suitable thresholds more complicated.

For this household the ranking of MaxOnly is poor in Fig. 9, because “black” days with nearly zero consumption are not consistently ranked as extremely likely holidays. While the ranking of days of logRegMax is plausible, too few days are detected as holidays because the threshold for p_d^i is set too low.




Fig. 9. Household 770: both MaxOnly (top) and LogRegMax (bottom) underestimate the number of holidays.

It is remarkable that even though households with a high number of holidays detected by MaxOnly have been selected for this section, still the number of holidays detected has often been too small. In order to demonstrate that such cases exist, as an example Fig. 10 shows the validation plots of household 952 where occupancy is obviously extremely rare. From a privacy perspective such households could be secondary residences that are only occupied for a small part of the year. Since in parts of Austria secondary residences are restricted finding such households as possible secondary residences could be a use case for government.

3.6 Households with Most Holidays for LogRegMax

Now households are analyzed, where logRegMax estimates more than 150 holidays but MaxOnly only estimates less than 100 holidays.

Households with 2 different characteristics occur. Households 55 and 892 have a high (where high here means near 1 kW) energy consumption for rather large parts of the day and also during parts of the night. While household 55 has no obvious holiday, household 892 has a clear holiday at the end of the measurement time. The validation plot of household 892 in Fig. 4 shows that both MaxOnly and logRegMax strongly overestimate the number of holidays. MaxOnly again fails because of the considerable energy consumption during night. LogRegMax




Fig. 10. Household 954: both MaxOnly (top) and LogRegMax (bottom) underestimate the number of holidays which is very large.




Fig. 11. Household 758: energy consumption shows vertical lines in winter which are possibly caused by an electric heater.

fails because the differences between the day and the night values are rather small. While the ranking makes more sense than for MaxOnly, the threshold is clearly set wrong resulting in an overestimation of the number of holidays.

The second set of households (758, 772, 940, 950) have the characteristic feature that vertical lines occur during winter time. As an example household 758 is shown in Fig. 11. Because of the primary occurrence in winter we suspect that these vertical lines stem from electrical heating systems that run throughout the day. Since it would make sense to turn the rather expensive electrical heating off during a holiday days with such vertical lines are likely not holidays.

Here, for both methods the bad distinction between night and day values lead to a failure of both methods as can be seen in Fig. 12. While the ranking of logRegMax is more plausible it still overestimates the number of holidays since many days with high consumption are left to the separation line.




Fig. 12. Household 758: MaxOnly (top) has a bad ranking of the days (top) and estimates 50, LogRegMax (bottom) overestimates 197 holidays.

Summarizing, in this regime, both methods suffer from overestimation for both kinds of household characteristics.

4 Validation for ECO Data

In order to build the validation on a stronger basis, another source of household consumption, the ECO dataset [9], is considered.

The heatmap of the 6 households is shown in Fig. 13. Although this data set is small, it poses a number of serious challenges for the holiday detection method. While houses 1 and 2 look typical and show longer periods of holidays, houses 3 and 4 have high average consumption throughout the year. Houses 3, 4 and 5 have a horizontal line of high values which is likely to stem from an automatic appliance. However, even though automatic appliances are not triggered by inhabitants, they can indicate occupancy. One example for such an appliance is automatic electrical heating. Houses 4 and 6 also show vertical lines that extend into the night. As seen before, using algorithms that rely too much on the distinction between day and night values will have problems there. Household 3 has an automatic appliance running during night between August and October. The behavior that it runs during night is expected to negatively affect algorithms that estimate thresholds from night values. The fact that it does not run through the whole year is an additional challenge, because the night values can be treated equally for all days. In the ideal case the regular structure of automatic appliances could be used to filter them out. However, such a method is currently not available. While the structure is rectangular it is not the same for all days which would certainly require a non-trivial removal method. The ECO data set illustrates a challenging diversity of load profiles, a detection should be able to deal with.

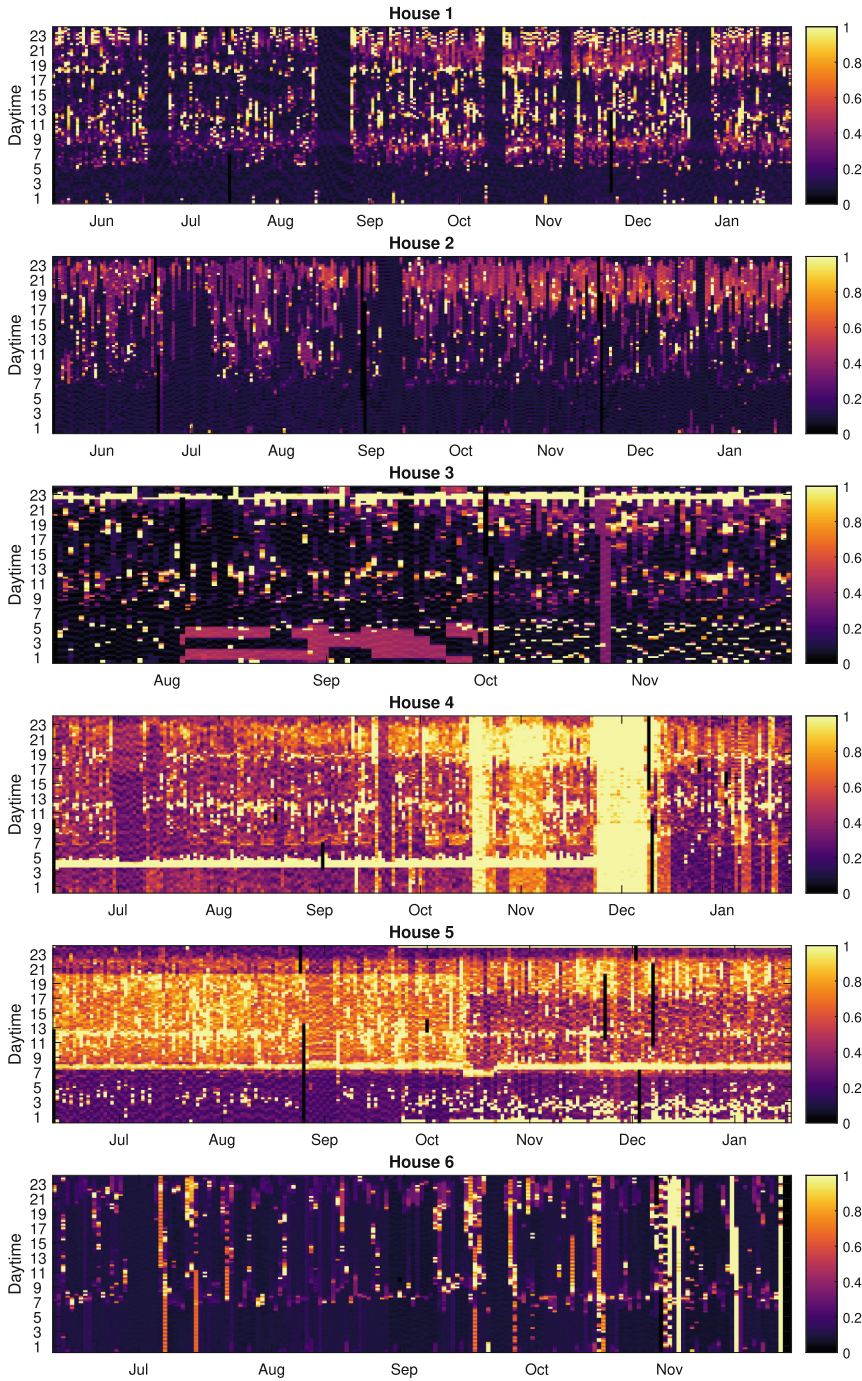


Fig. 13. Heatmaps describing the 6 houses of the ECO dataset.

Both MaxOnly and LogRegMax were applied to all households. The validation plots are shown in Figs. 14 and 15, respectively. MaxOnly showed plausible results for houses 1, 2 and 6 with a slight overestimation again for days with high night values (leftmost part of the plot). The overestimation was moderate for household 3, it was extreme for households 4 and 5. LogRegMax showed a plausible ranking for households 1, 2, 3 and 6 with underestimation for households 1, 2, and 6. As MaxOnly, it extremely overestimated the number of holidays for households 4 and 5.

In summary, both methods showed big weaknesses for the ECO data set.

5 Towards a New Model

The validation of the existing methods in Sects. 3 and 4 showed that there is much for improvement. It also showed that the problem is not as simple as initially considered. MaxOnly generally suffers from false positives when higher values occur during night. LogRegMax mainly suffered from the setting of the threshold especially when most of the days are holidays or when consistently high values also occur during night.

5.1 Use Only Maximum Values: MaxTol

The methods proposed so far suffered when night and day values were too similar. This characteristic leads to a big error because the number of holidays are overestimated. While it has been clear that this is an explicit assumption naturally arising from the adaption of the occupancy detection approach in [2], the experiments showed that the distinction between day and night may introduce more problems than it solves. Therefore, here we propose to skip the distinction between night and day values.

Another way to treat this problem would be to try to remove automatic appliances that also run during night time thus decreasing the difference between daytime and nighttime consumption. However, such methods are not yet readily available for all kinds of appliances, especially for low resolution data. Therefore, we leave this way as a topic for future research.

Especially houses 4 and 5 of the ECO data showed a clear overestimation. For a human, occupancy seems to be clear that somebody is at home, if a “high” load value occurs. In contrast to this intuition, the current methods all only compared day values with night values, but no method compared the *absolute* value with a threshold. Therefore it seems reasonable to consider a fixed threshold for the maximum of the whole day. Using such a threshold is supported by the results about the determinants of electricity consumption [6], where it is stated that “most high consumption, intermittent appliances such as electric water heater, electric clothes dryer, and Spas/ Pools primarily contribute to daily maximum consumption. These are the appliances that are not “always on” and their operating schedules are dependent on the activities and habits of the occupants”. The power consumption of the class of ohmic devices that need

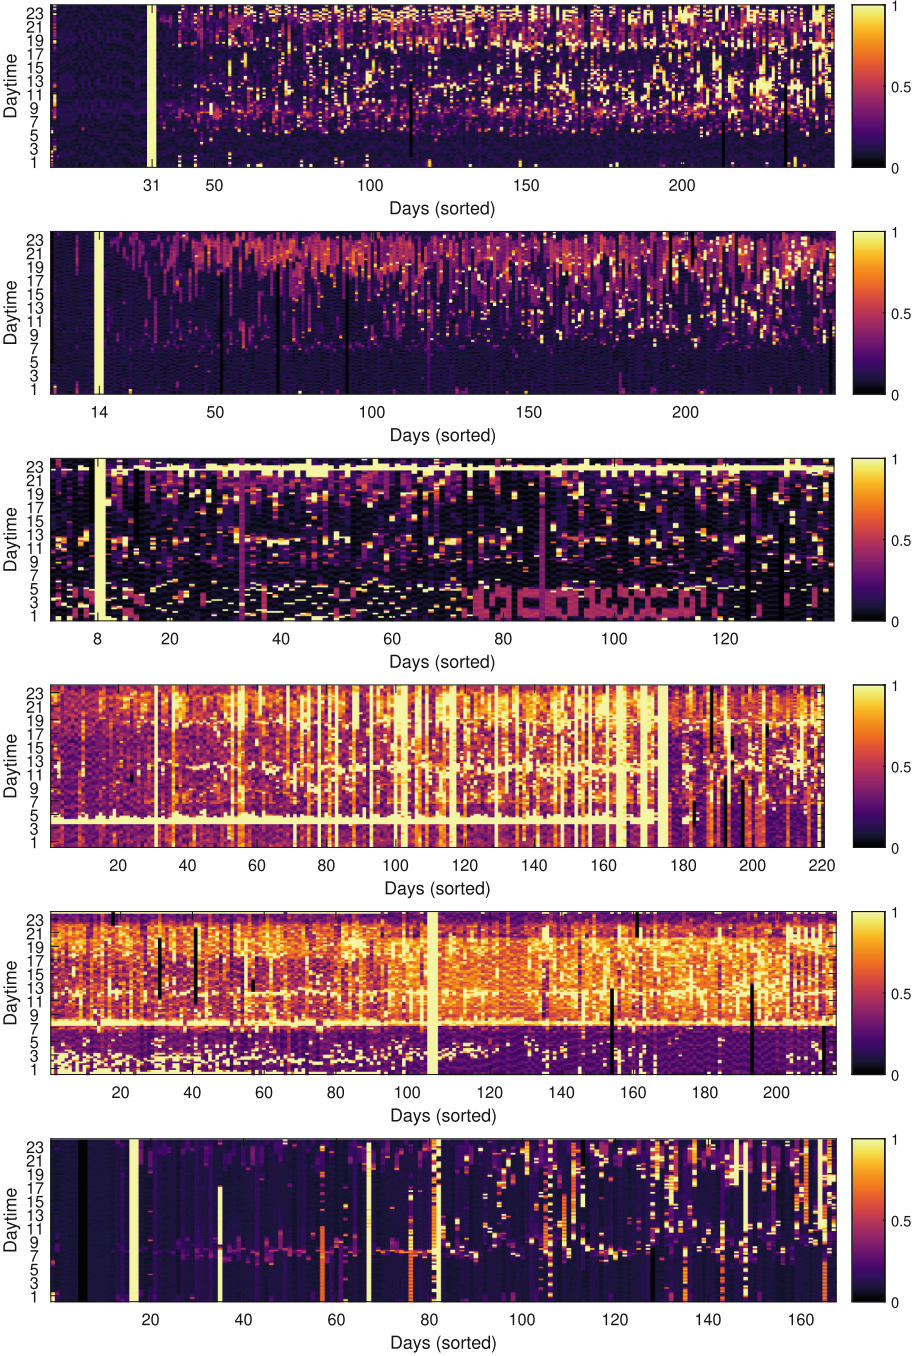


Fig. 14. Validation of MaxOnly for ECO data.

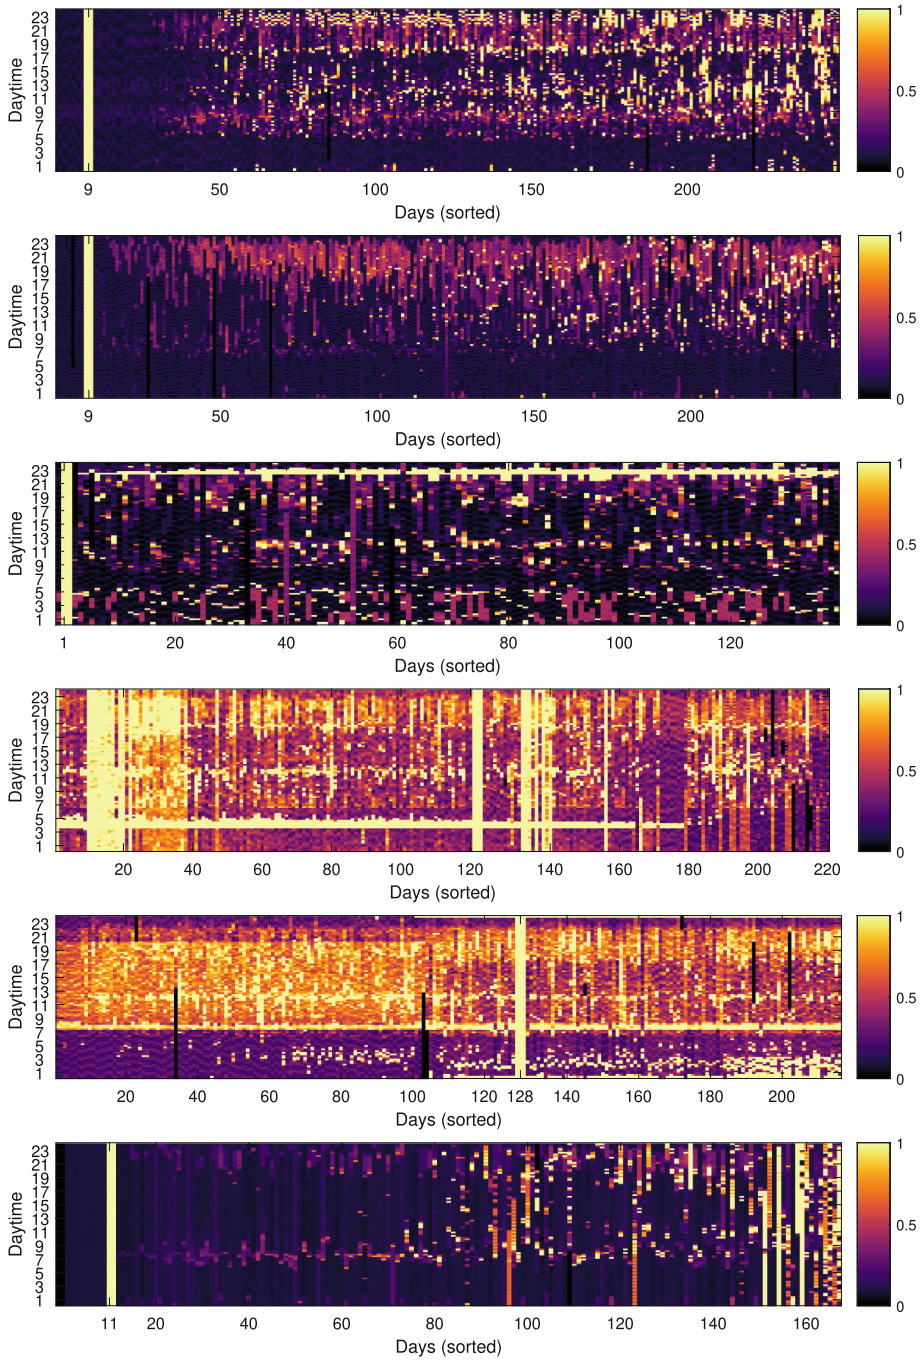


Fig. 15. Validation of LogRegMax for ECO data.

to produce heat have power values above 1 kW. Therefore, we used 1 kW as an upper limit for a value where we can be quite sure that it stems from an activity of a human in the household. Consequently, all values entering a heatmap have been clipped from above at 1 kW so that the heatmaps are suitable to show the low maxima that are interesting for developing thresholds. The fact that this clipping enabled us to conduct the validation analysis based on the heatmaps supports the usefulness of considering values above 1 kW as likely human-driven.

Using these considerations as a guideline, we next study the following, extremely simple rule-based classifier that compares the maximum value of the *whole* day with a constant, household-dependent threshold. Consider a whole day d of household i . Denoting the 96 quarterly-hour power values of as $\mathbf{x}_{i,d} = (x_{i,d}^1, \dots, x_{i,d}^{96})$, the binary holiday variable $h_{i,d}$ is assigned by the following simple rule using a single tolerance value Tol_i for a household

$$h_{i,d} = \begin{cases} 1 & \text{if } \max_{t \in \{1, \dots, 96\}} \mathbf{x}_{i,d}^t < \text{Tol}_i \\ 0 & \text{if } \max_{t \in \{1, \dots, 96\}} \mathbf{x}_{i,d}^t \geq \text{Tol}_i \end{cases} \quad (2)$$

For future reference, this simple rule is called MaxTol. From (2) it is clear that more holidays will result from choosing a bigger value of Tol_i .

Two questions immediately arise. First: is such a simple rule flexible enough to get a reasonable holiday detection mechanism? Second: how should the threshold be determined? For these 2 questions preliminary answers will be given in the next two subsections based on trials for the diverse ECO data set.

5.2 Threshold Selection

Before demonstrating how the threshold is chosen, the effect of the threshold on the outcome is illustrated in Fig. 16 for first, bigger dataset. There, the threshold has been set to the same value for all households.




Fig. 16. Effect of different thresholds on the holiday distribution.

From Fig. 16 two conclusions can be drawn: first, the threshold has a huge effect on the distribution of the number of holidays. Second, the range between

0.2 and 1 suffices for realistic, typical households. If Tol would be set to 1 for all households, around 30% of households would have 80 holidays or more per year. This seems unrealistically high for Austria.

Threshold selection is done manually for each household based on the heatmap of a household and the cumulative distribution of the night values (Fig. 17). In the first step, the night time is determined based on the heatmap. For all households of the ECO dataset the night time was set to the time span between 2 h and 4 h, only for house 4 the time span between 0 h and 2 h was selected in order to avoid the values of the “horizontal line”. Then the cumulative plot of the night values was created (Fig. 17).




Fig. 17. How threshold have been selected (red, dashed line) based on the cumulative distribution of the night values (blue). (Color figure online)

This plot serves as a basis to distinguish high values that indicate activities stemming from humans from the background of automatic appliances. Based on inspection of the heatmap (Fig. 13), if no high values occur during night, the threshold was set high enough to include nearly 100% of the night values (houses 1 and 2). If a considerable amount of high values exist, the threshold should be chosen to sort these values out. For example, house 3 contains a considerable amount of high values. The plot of the cumulative distribution shows that these high values account for about 30% of the night values that can be seen in the

plot as the second increase after the plateau. The threshold (red line) was set to just exclude these values. In a similar way the threshold can be achieved for houses 4 to 6. Only for household 4 the position of the threshold stays rather unclear since no clear flat intermediate region occurs. Based on several manual trials it turned out that for house 4 the result was insensitive to the exact setting of the threshold.

As already can be seen in Fig. 17 the thresholds are well below 1 kW but changed considerably between houses. The thresholds used are also listed in Table 3.

Table 3. Results of MaxTol for ECO data.

House ID	1	2	3	4	5	6
Used thresholds Tol_i	0.2	0.2295	0.2395	0.496	0.5	0.13
Detected number of holidays	27	11	0	0	0	11

5.3 Application to ECO Data

Finally, using the thresholds derived in Sect. 5.2, MaxTol was applied to the ECO dataset. The corresponding validation plot in Figs. 18 and 19 ordered the days by increasing maximum values, now using values from the *whole* day for the determination of the maximum.

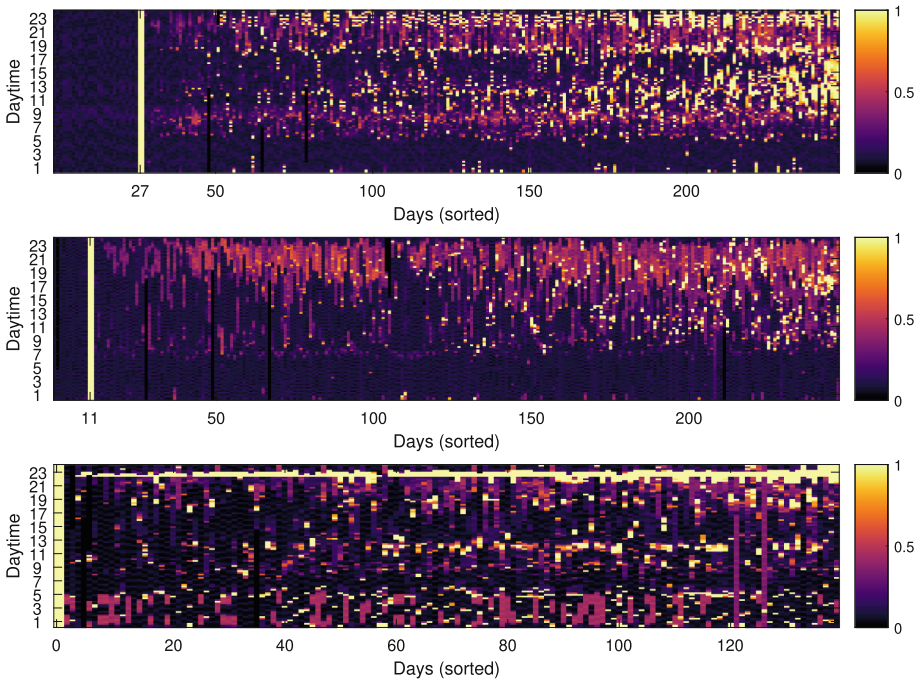


Fig. 18. Application of MaxTol on ECO data houses 1, 2 and 3.

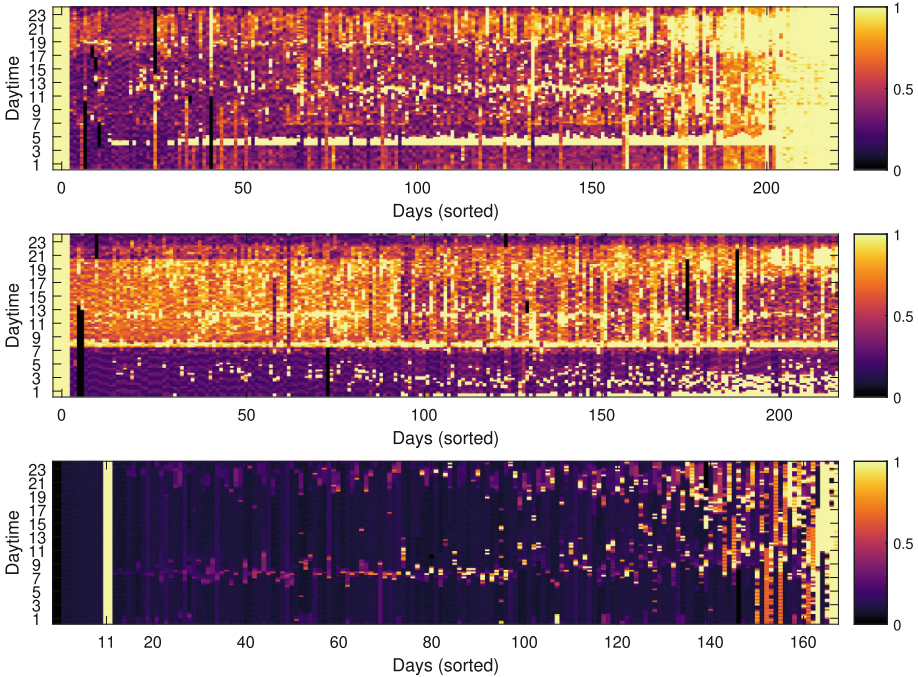


Fig. 19. Application of MaxTol on ECO data houses 4, 5 and 6.

The figures show that reasonable results can be achieved by MaxTol for all households although it is in principle simpler than the other investigated methods, which performed weak for this dataset. However, this good result comes at the cost of needed manual intervention. It should be noted that a few selected households from the other data set have also been successfully treated this way.

6 Conclusion and Outlook

In this paper, previously developed methods that detect holidays by analyzing low frequency smart metering energy consumption data have been evaluated. It turned out that these methods work well for normal cases but have big weaknesses for various situations. The analysis revealed several issues that must be considered: consistently high values for both night and day, automatic appliances that run for parts of the day or night but not all year, background that depends on the time of the year or also households with unexpectedly rare occupancy over the year. From a privacy perspective, it can be considered good news that holiday detection is not as simple in general as one would think at first.

The analysis showed that the automatic distinction between day and night values posed more problems than it solved. As a first step towards an improved method we therefore propose to use the maximum value of the *whole* day.

There is also evidence that it is better to only use the maximum without additional quantiles of the load values. The absolute value of thresholds should be restricted to be in a certain range that is below consumption values of known human-driven appliances like ohmic heaters. Based on these principles a more suitable, very simple method is proposed. It is shown that using this simple method in principle the houses of the ECO dataset can be treated in a plausible way. However, the method comes at the cost of human intervention in estimating the threshold.

Holiday detection from electric load profiles is still in its infancy. Therefore many extensions are possible. A particular challenge will be to find thresholds automatically. It could also be challenging to develop more flexible models that are based on modeling of human habits and activities. Validation currently needs visual inspection which is both laborious and error-prone. Because labeled data will be hard to achieve, validation by household simulations could be a viable alternative.

Acknowledgement. The financial support by the Federal State of Salzburg is gratefully acknowledged. Furthermore, the authors would like to thank the Energieinstitut at the Johannes Kepler University Linz for providing the data set.

References

1. Becker, V., Kleiminger, W.: Exploring zero-training algorithms for occupancy detection based on smart meter measurements. *Comput. Sci. Res. Dev* **33**(1–2), 25–36 (2018). <https://doi.org/10.1007/s00450-017-0344-9>
2. Chen, D., Barker, S., Subbaswamy, A., Irwin, D., Shenoy, P.: Non-intrusive occupancy monitoring using smart meters. In: *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings - BuildSys 2013*, pp. 1–8 (2013). <https://doi.org/10.1145/2528282.2528294>
3. Eibl, G., Burkhart, S., Engel, D.: Unsupervised holiday detection from Low-resolution smart metering data. In: *2018 Proceedings of the 4th International Conference on Information Systems Security and Privacy, ICISSP*, pp. 477–486. SciTePress (2018). <https://doi.org/10.5220/0006719704770486>
4. Hart, G.W.: Nonintrusive appliance load monitoring. *Proc. IEEE* **80**(12), 1870–1891 (1992)
5. Jin, M., Jia, R., Spanos, C.: Virtual occupancy sensing: using smart meters to indicate your presence. *IEEE Trans. Mob. Comput.* **16**(11), 3264–3277 (2017). <https://doi.org/10.1109/TMC.2017.2684806>. <http://ieeexplore.ieee.org/document/7882676/>
6. Kavousian, A., Rajagopal, R., Fischer, M.: Determinants of residential electricity consumption: using smart meter data to examine the effect of climate, building characteristics, appliance stock, and occupants' behavior. *Energy* **55**, 184–194 (2013). <https://doi.org/10.1016/j.energy.2013.03.086>
7. Kim, H., Marwah, M., Arlitt, M.F., Lyon, G., Han, J.: Unsupervised disaggregation of low frequency power measurements. In: *The 11th SIAM International Conference on Data Mining*, pp. 747–758 (2011)

8. Kleiminger, W., Beckel, C., Santini, S.: Household occupancy monitoring using electricity meters. In: Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing, pp. 975–986 (2015). <https://doi.org/10.1145/2750858.2807538>
9. Kleiminger, W., Beckel, C., Staake, T., Santini, S.: Occupancy detection from electricity consumption data. In: Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings - BuildSys 2013, pp. 1–8 (2013). <https://doi.org/10.1145/2528282.2528295>, <http://dl.acm.org/citation.cfm?doid=2528282.2528295>
10. Lisovich, M.A., Wicker, S.B.: Privacy concerns in upcoming residential and commercial demand-response systems. In: Clemson Power Systems Conference. IEEE (2008)
11. Zoha, A., Gluhak, A., Imran, M.A., Rajasegarar, S.: Non-intrusive load monitoring approaches for disaggregated energy sensing: a survey. *Sensors (Switzerland)* **12**(12), 16838–16866 (2012). <https://doi.org/10.3390/s121216838>