Black Box Approach for Energy Monitoring of Commercial Buildings

Susanne Komhard
M.Sc. student
Fraunhofer-Institute for Solar Energy Systems
Freiburg, Germany

Christian Neumann
Project leader
Fraunhofer-Institute for Solar Energy Systems
Freiburg, Germany

ABSTRACT

The potential to save energy by changing operational parameters - especially in existing commercial buildings – is in the magnitude of 5-30%. In order to realize this saving potential in the long term, continuous commissioning of the building is a key issue. Necessary for successful continuous commissioning is real time monitoring of the building performance which allows for Fault Detection and Diagnosis (FDD). This paper presents a method to monitor building operation and detect faulty or unusual behaviour using a black box model approach. The approach is to identify a building's basic operating characteristics by means of measured data from a building to train a multiple linear regression model based on energy signatures of the building. In addition to supplying measured building data to the regression a clustering process is added which determines the building's day-types. Once the model is trained it can predict the energy consumption at the building site and unusual or faulty days can be identified by comparing the predictions to real measurements. Models to monitor the daily heating and electricity demand are developed and applied to measured data from two demonstration buildings.

Key words: multiple linear regression, performance monitoring, commercial buildings, Black Box Models, FDD, outlier detection, clustering, day-types, energy signature

INTRODUCTION

The idea to use multiple linear regression to identify a building's operating characteristics originates from energy signatures which display the linear relationship between energy demand and outdoor air temperature, as visible in Figure 1.

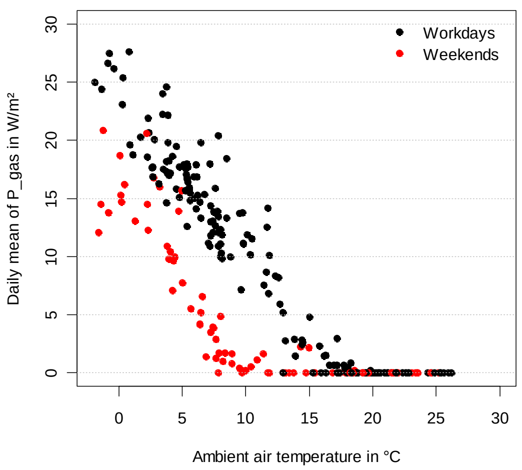

Figure 1: Energy signature of heating demand

In the case of heating demand this relationship is linearly decreasing and is mathematically represented by a simple linear equation with negative slope and intercept. In multiple linear regression it is possible to use more explanatory variables than merely the outdoor air temperature to explain an energy demand. In the building context this means that also the indoor air temperature, water consumption, electricity consumption and further measurement points can be integrated into the regression to achieve a well fitted model. The multiple regression serves as a black box model. If supplied with enough measurements it can estimate the unknown parameters of slope and intercept.

Visible in Figure 1 are four different operating characteristics which can be distinguished by using day-types and a change point. Most buildings have two different day-types which describe the energy consumption during workdays (black) and weekends (red). The cases have similar slopes but different intercepts. Change points are given in degrees of Celsius of outdoor air temperature and describe the temperature at which a change of operation occurs, namely the transition from heating or cooling demand to base load level. The different operating characteristics lie either above or below the change point and have distinct slopes and intercepts.

The identification of day-types is carried out by applying an agglomerative hierarchical clustering process which distinguishes between days that show similar energy consumption profiles and days that do not. The idea to use clustering to identify day-types is derived from Seem [1].

The information of day-types and change point is made available to the multiple regression in form of categorical variables which are either true or false, depending on the day.

MODEL OVERVIEW

The model is separated into two main parts: Training and Validation visible in Figure 2 and Figure 3. The Training part is itself separated into six stages starting with *T1* which compiles 15 minute or one hour measurements to daily averages in W/m² of the building's net floor area or in °C. The starting point is to use at least three months of data to provide sufficient measurement points to the regression.

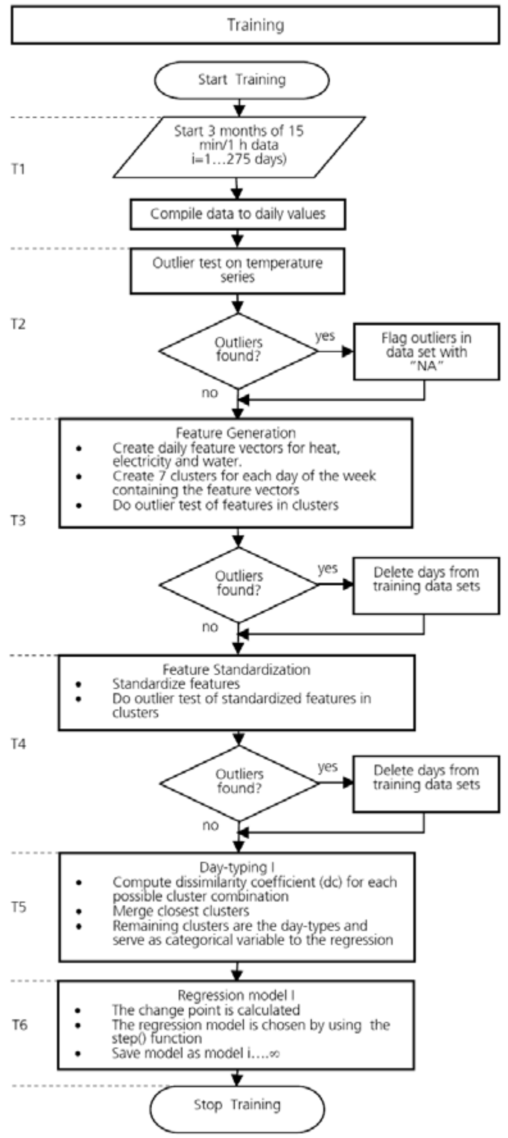

Figure 2: Training stages of model

Step *T2* searches temperature series such as outdoor and indoor air temperatures for measurement errors The next section *T3* generates six daily features for the clustering process. Features are derived from heat, electricity and gas consumption and represent the average daily and maximum hourly consumption in one day. The six features serve as the building's daily consumption characteristics for the clustering process. To be able to exclude measurement errors and unusual days, the features are searched for outliers prior to their standardization. Standardizing the features is step *T4*. The applied method is explained further down. Also part of step *T4* is another outlier detection method to remove unusual days which come forward due to standardization. After the features are created the clustering process determines the building's day-types which is step *T5*. Its result is used for the multiple regression model in form of a categorical variable. The last step is the regression model which is trained by using the daily averages created via the data compilation (*T1*), the information about the day-types (*T5*) and the temperature value of the change point. The change point is the result of optimizing the regression model in regard to the adjusted R^2 value of the regression. Once the regression is trained the Validation part begins in which the model can be used to predict the building's daily energy demand when provided with the necessary explanatory variables such as temperatures, water consumption and energy demand other than the target variable's. The prediction is thus not a forecast but carried out in retrospect for the energy demand of a preceding day. The predicted value is compared to the actual demand and days with unusual consumption can be identified. This is the Validation stage presented in Figure 3.

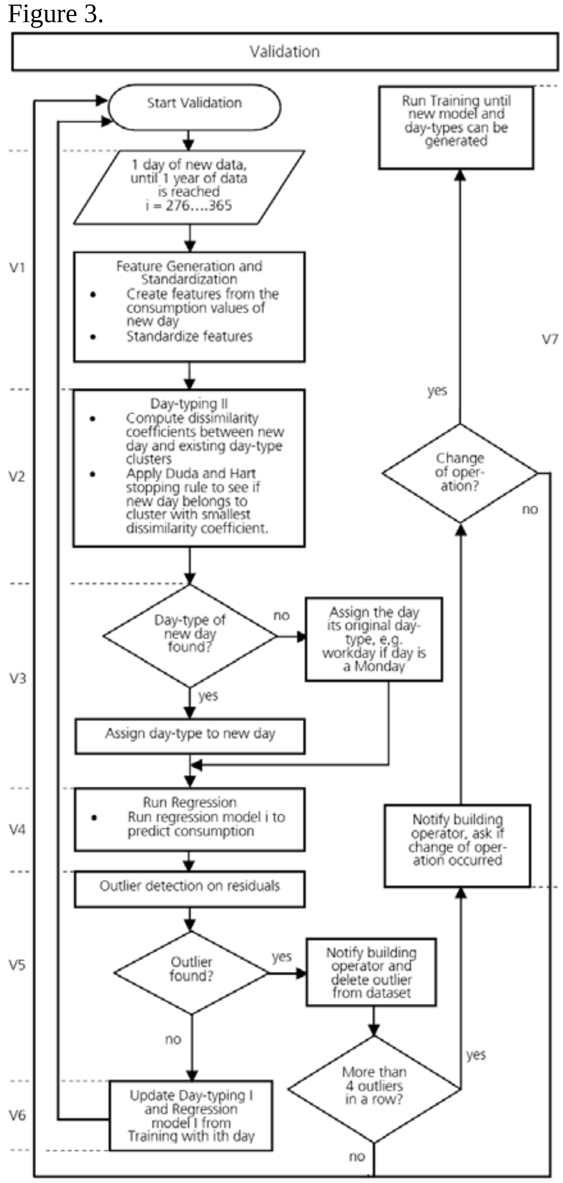

Figure 3: Validation stages of model

Validation starts with *V1* and describes the process of generating and standardizing the six features of a new day. The features are derived from the day's heat, electricity and water consumption, identical to the Training part. The next step, *V3*, is to test whether the day belongs to an existing day-type from the Training stage. If this is not the case the day is assigned its original day-type. Otherwise the day is assigned the found day-type and the regression model is used to predict the day's energy consumption which is step *V4*. If the prediction deviates considerably from the measured value the building does not operate as usual and the day is an outlier. Testing whether a day is an outlier is marked with *V5*. Only sound days are added to the data from the training stage with which regression model and clustering process are updated. If more than four days are unusual in a row it can be assumed that a change of operation in form of additional HVAC-appliances has taken place. In such a case regression and clustering have to be trained anew with sufficient data representing the new building operation. The next sections present in detail the different stages such as feature generation and regression model.

FEATURE GENERATION AND CLUSTERING

For each day during the training period six features are generated: The average daily and the maximum hourly consumption of electricity, heat and water. It became apparent that using these six features to characterize one day offered good results. Adding also the minimum hourly consumption can result in features that contain a zero value every day. Such features do not comprise any useful information to describe different days because they contain the same value and are not distinct from each other. Once the features for each day are generated they are put into eight sets of predefined clusters which are: Mondays, Tuesdays, Wednesdays, Thursdays, Fridays, Saturdays, Sundays and Holidays. After the features are put into their respective cluster they are searched for outliers. This process should catch features with extreme values which are, first, due to erroneous measurements and, second, describe days that are different from the majority in their cluster. For this purpose three times the standard deviation is taken as outlier detection limit. If any feature contains such a large deviation the whole day is assumed to be erroneous and it is deleted from its cluster and from the training set containing the mean daily values for the regression model.

The approach to use clustering for the determination of day-types has previously been carried out by Seem [1]. Seem uses two features to characterize days which are then transformed to account for changes in consumption patterns. Those changes are mainly due to seasonal influences such as high gas consumption in winter when heating is switched on and low consumption in summer when gas is only needed for hot water production. In his paper he transforms the features by subtracting from each daily feature 1/7th of the value of the three subsequent, three preceding features and the feature itself. Applying this procedure to the measurements of the demonstration buildings it showed however, that seasonal influences were not satisfyingly removed and a different approach was applied: During the training stage, a moving window the size of seven days is shifted over every week in the 90 day long training period. This results in approximately 13 windows. Each window selects the maximum value of the six features during the seven day period and divides the features by their respective maximum. Like this, a feature which is a maximum during winter times has the value one and a feature which is a maximum during summer times has a value of one, too. Seasonal changes are removed completely without loosing any information on the consumption profile. Additionally, the features are standardized and always range between zero and one. The following figure shows a comparison between original features, transformed features according to Seem and the features standardized according to the method applied here:

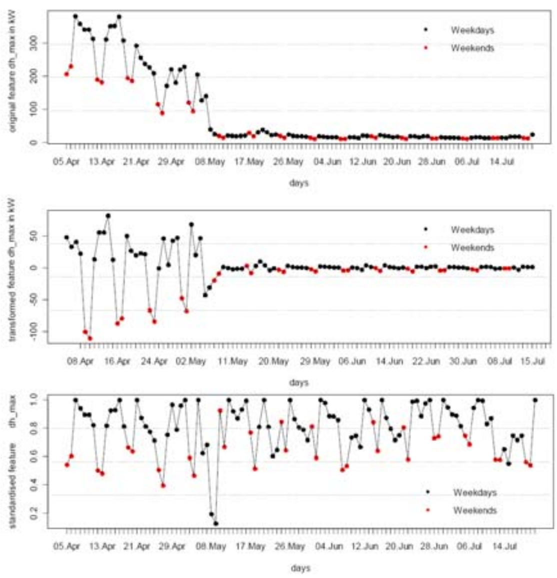

Figure 4: Original, transformed and standardized features of district heat consumption

The top graph of Figure 4 shows the original feature of the maximum hourly demand during one day for the district heat consumption at one of the demonstration buildings. After the heating has been switched off the profile passes over to base load level with the result that the magnitudes of weekend and workday consumption are almost identical. Applying the standardization according to Seem, visible in the middle graph, does not change this profile. But by using the feature standardization visible in the bottom graph of Figure 4, the differences in the consumption profile between weekends and workdays are now also available during the summer time. The extreme gap for the two days around May 8th is due to a sudden drop in consumption when heating is switched off. The features indicate the first two days when no heating takes place. But they fall into the seven day window when heating is still running and are thus divided by a comparatively very large maximum value. If the time series started two days earlier, the two features would fall into a window containing only features representing the stage when heating is switched off and the standardized features would be in the normal range. But as such sudden changes in operation are likely to happen it is necessary to run a second outlier test on the standardized features. A day is fully deleted from the feature's training data set if two or more standardized features contain unusual values. The reason for applying another outlier test to the standardized features is also apparent in following figure:

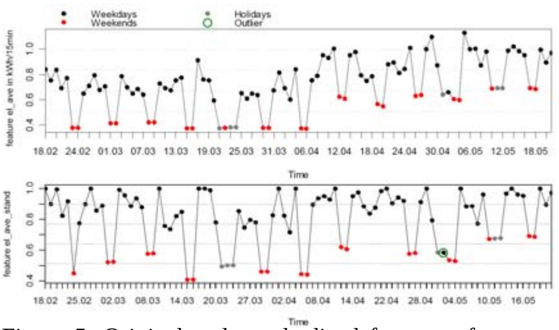

Figure 5: Original and standardized features of average daily electricity consumption.

The upper graph shows the feature of the average daily electricity consumption in original and standardized form. One particular day is unusual marked with a green circle. The day is Friday the 2nd of May and is located between a public holiday and a weekend. Most people took this day off work and the consumption profile is unusually low for a Friday. The value should not be part of the Friday cluster. The timeframe chosen contains a change of operation in which multi-split units are switched on indicated by the increase in consumption in the upper graph. The reason why the unusual Friday is not found in the first outlier detection run is that the un-standardized feature el_{ave} contains low (spit units off) and high (split units on) values on Fridays and the 2nd of May is no outlier. Due to standardization the change of operation of adding split units is removed and the feature on the 2nd of May is now lower than on other Fridays and can be detected by the outlier test.

It is just as important for the clustering process as it is for the regression to remove days like the 2nd of May. Keeping such unusual days in the training data set causes a shift in the regression plane with the result that predictions for workdays contain slightly lower values.

The next step after putting the features in their respective cluster, standardizing them and testing them for outliers is to evaluate how close the clusters are situated to each other. The process is identical to the one Seem applied [1]. The first step is compute dissimilarity coefficients based on the Euclidean distance between two clusters. The two clusters which are situated the closest have the smallest dissimilarity coefficient. These clusters are evaluated according to the stopping rule of Duda and Hart [2] which compares the sum of squared error of the clusters when they are separated and when they are merged. Clusters which do not satisfy the stopping rule can be merged.

For the majority of office buildings clustering results in two distinct day-types: Mondays, Tuesdays, Wednesdays, Thursdays and Fridays which form the "workday" day-type and Saturdays, Sundays and Holidays which form the "weekend" day-type.

MULTIPLE LINEAR REGRESSION

The regression model is supplied with the daily averages of the training period, the different day-types and a change point to estimate the regression coefficients (slope and intercept). A regression model for the heating demand has following form:

$$P_{heat} = \beta_0 + \beta_1 (T_a - CP) + \beta_2 P_{el} + \beta_4 V_w + \beta_5 T_i + \beta_6 \Delta T_i \quad (1)$$

Where

P_{heat}	daily mean of specific heating demand in W/m² of net floor area
β_i	parameters of the model
T_a	daily mean of outdoor air temperature in °C
P_{el}	daily mean of specific heating demand in W/m² of net floor area
V_w	daily mean of water consumption in l/h
T_i	daily mean of indoor air temperature in °C
ΔT_i	difference of daily mean of indoor air temperature to the previous day in °C
CP	change point

Eq. (1) is calculated for all levels of day-types and above or below the change point. Assuming two different day-types and one change point the model creates four different cases. Each case has different intercepts and different slopes.

An additional calculated variable ΔT_i is the temperature difference between two subsequent days. The variable accounts for an increased cooling down or warming up of the building envelope whenever cooling or heating has been set back. This is mostly the case during weekends which in return is followed by an increased heating or cooling demand on the subsequent day. Adding ΔT_i provides the necessary information to the regression regarding cooling down and warming up processes in the building.

The result of clustering and regression is on the one hand side the building's characterization according to day-types and on the other hand a regression model that describes the building's energy demand at different operating conditions. However, a regression model which is trained from data that only represents the operating conditions during winter is not able to make predictions for summer times. Thus, it is advantageous to supply training data which represents different operating characteristics which mostly coincides in spring or autumn.

FIELD TEST RESULTS

The above described methods of Training and Validation are tested with measured data from two demonstration buildings. The results are presented in graphs which show a comparison of measured and predicted values from the regression.

The first model is developed for the heating demand. A three month training period is chosen. The result is visible in Figure 7. The graph is to illustrate how the trained regression model fits to the actual measurements. Data points marked with a green circle are not used in the training of the model but the model predicts their values when given the corresponding explanatory variables.

The model has been trained with three day-types which were the result of the clustering process for the available training period. The day-types are “workdays”, “weekends” and “holidays”. Due to the fact that the holidays in the Training period show a consumption profile which is similar neither to weekends nor workdays they were not merged to any other day-type and form their own cluster.

The regression model has a very high adjusted R^2 value of 0.96 which means that the explanatory variables can explain 96% of the variation of the heating demand. To make sure that this explanatory power is provided by all explanatory variables the regression model is searched for statistically insignificant variables which are excluded from the model. In the case of the heating demand all available variables according to eq. (1) are significant and are part of the model. The change point of the model is at 13 °C. Judging from the high quality of the model it should be well capable to predict correct values in the Validation stage.

The Validation stage is chosen to be before the

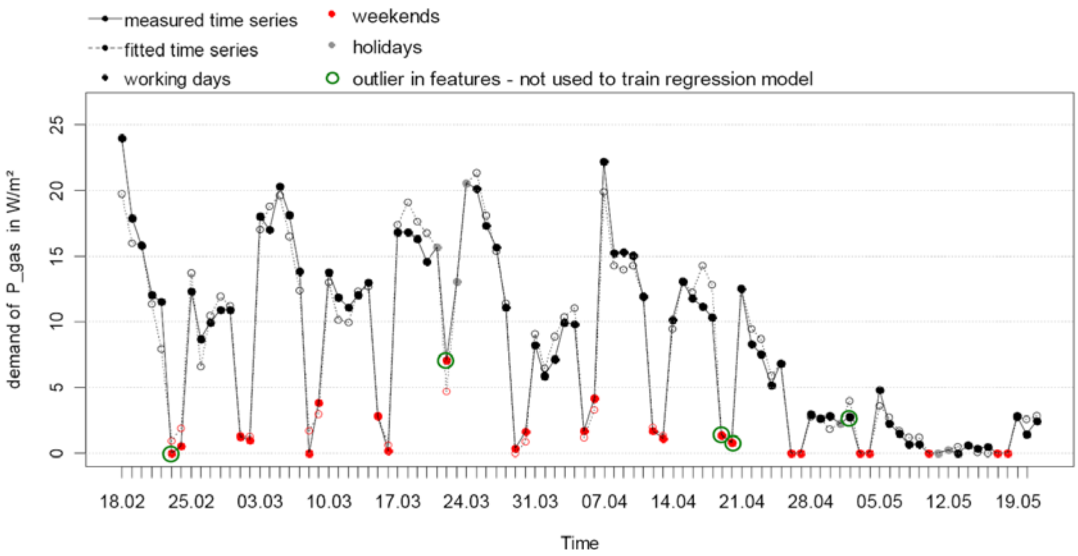

Figure 7: Result of Training for heating demand showing fitted values and measured values.

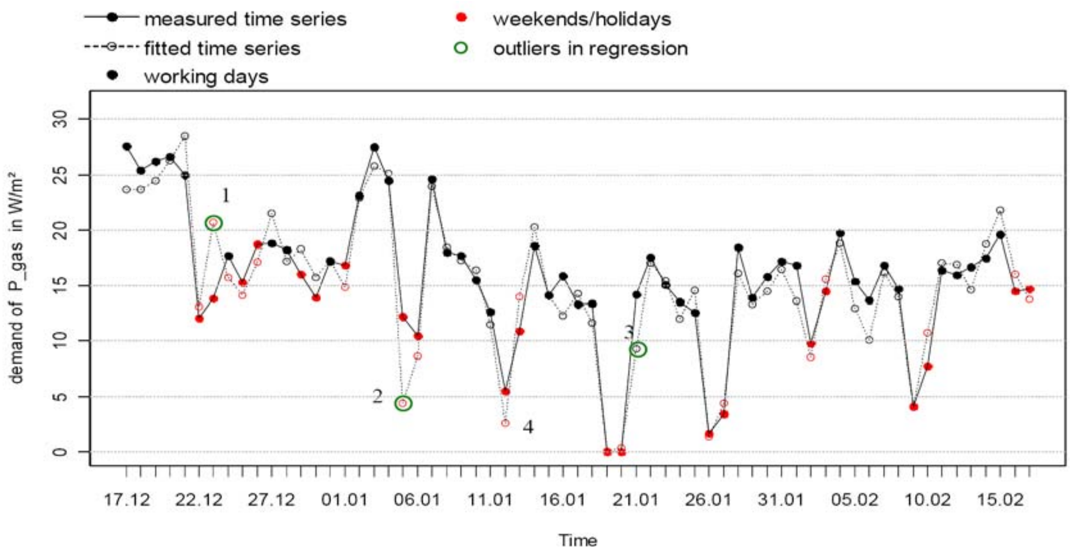

Figure 6: Result of Validation for heating demand showing predicted values and measured values.

Training timeframe so that it is possible to test the models efficiency during times of heating demand. The result is presented in Figure 6. Three outliers are found during the Validation stage. The first is on the 23rd of December. The prediction is almost 7 W/m² higher than the actual measured value. One of the reasons why the validation period is chosen to be during Christmas time is that it appears odd that the gas demand is comparatively high. But looking at the predictions from the model during this time it seems that the gas demand is reasonable. The reason for this might also be in the very low outdoor air temperatures during this time which demand higher gas consumption. The 24th of December is assigned the day-type "weekend" which shows that the clustering process correctly evaluates the building's operating situation because even though the 24th of December 2007 was a Monday people were not working. The second day which is an outlier is the 5th of January. The assigned day-type "weekend" is correct and the prediction is almost 8 W/m² below the measured value. It seems that the weekend set back of the heating has not been applied as efficiently as on other weekends.

The third day is the 21st of January when the prediction is around 5 W/m² higher.

The Saturday which is numbered with a four also shows higher gas consumption than what the model predicted. The residual is just slightly below the limit to be marked as an outlier.

To further test the reliability of clustering and regression two outliers are manually added to the 15 minute measurements of the gas consumption. Adding them to the 15 minute measurements ensures that they are contained in the features of the clustering as well. The first outlier increases all measurements of the gas consumption on Tuesday the 22nd of January by 20 %. This outlier should be easily caught by the regression model. The second outlier is added to Saturday the 26th of January by substituting the weekend gas consumption with the profile of a workday. The values of P_{gas} from the previous day, Friday the 25th, are copied into the 15 minute measurements of Saturday the 26th. Here it is important that the clustering process still recognizes the day type "weekend" even though the features of the gas consumption say different. But the features of the water and electricity consumption still indicate "weekend" and should be sufficient for a correct day-type determination. The results are plotted in Figure 8. As visible both outliers are caught by the regression which predicts lower values. Especially nice is the fact that the clustering recognized the outlier's day-type correctly and that the prediction for Saturday is made with the weekend case. If the clustering process had assigned the day-type "workday" the regression's prediction would be higher and the day would not be an outlier.

The example concerning the heat demand shows that the clustering process works very well. Days which are weekends are put into the "weekend" cluster and workdays are put into the "workday" cluster. The regression models predictions are also very close to the actual measured values. The outliers found are realistic. The days numbered 1 and 4 in Figure 6 are a good example of the potential of the model algorithm: Days on which more energy than normally is used have a lower prediction and can be separated from days which show a normal building operation. The reliability of clustering and regression model is also confirmed by the manually added outliers which were correctly identified.

The second example is presented in form of a model to monitor the electricity demand. The electricity measurements are from a demonstration building that mainly uses electricity for plug loads and about 10% are used for running multi-split

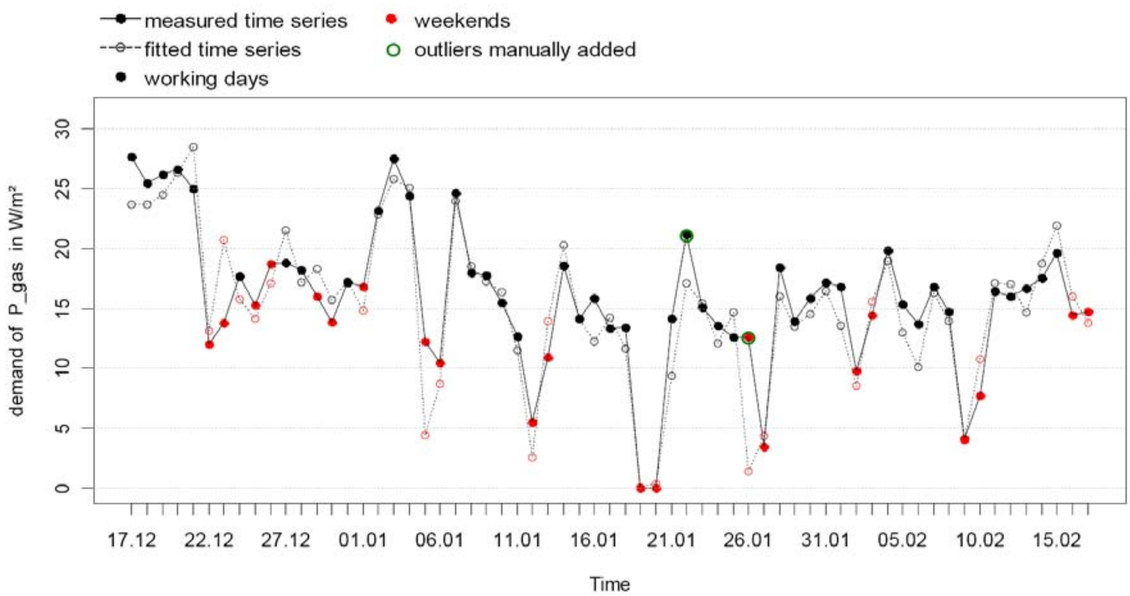

Figure 8: Result of Validation for heating demand showing predicted values and measured values with added outlier days.

units. The result of the training stage is presented in Figure 10. It shows that the building's consumption profile on Saturdays is slightly different to Sundays. The clustering process should result in three different day-types: "workdays", "Sundays" and "Saturdays". To achieve this it was necessary to keep the features in their un-standardized form. If this is merely a coincidence or due to the fact that too much information regarding the scale and location of the measurement points is lost during standardization is not clear and further investigations into that matter will be carried out. Keeping Saturday as a separate day-type is very beneficial for the regression. It shows that the model fit for weekends is almost perfect. The model has an adjusted R^2 value of 0.95.

The most important days to be removed from the training data set are the two Fridays which show a lower consumption profile than all other weekdays. The predictions of the trained regression model on these days are a lot higher. If the two days were used to train a model its predictions would be closer to the measured values with the result that the regression coefficients would have to be different influencing all other predictions, as well.

The Validation period is shown in Figure 10. The predicted values are very close to the measured values. Except for the 25th of July which contains a higher prediction no other days are found to be outliers. Days are outliers when the predicted value is more than 15 % off the highest measured value. The difference between the predicted and the measured value on July the 25th is just slightly above this limit.

The regression model for the electricity demand is different to the one for the heating demand because it does not contain the explanatory variables for the indoor air temperature T_i and the temperature

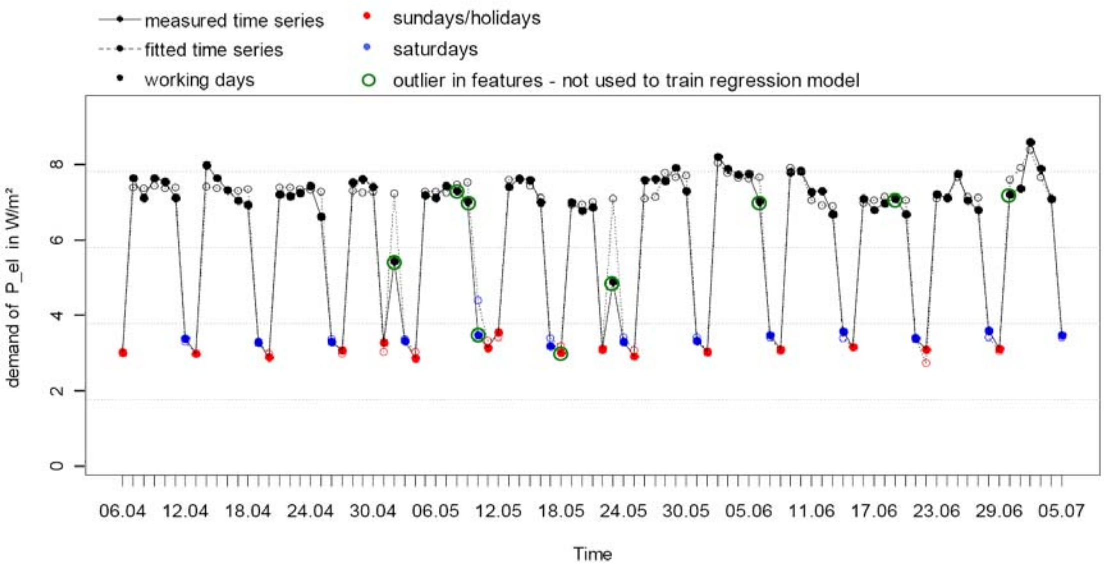

Figure 10: Result of Training for electricity demand showing predicted values and measured values.

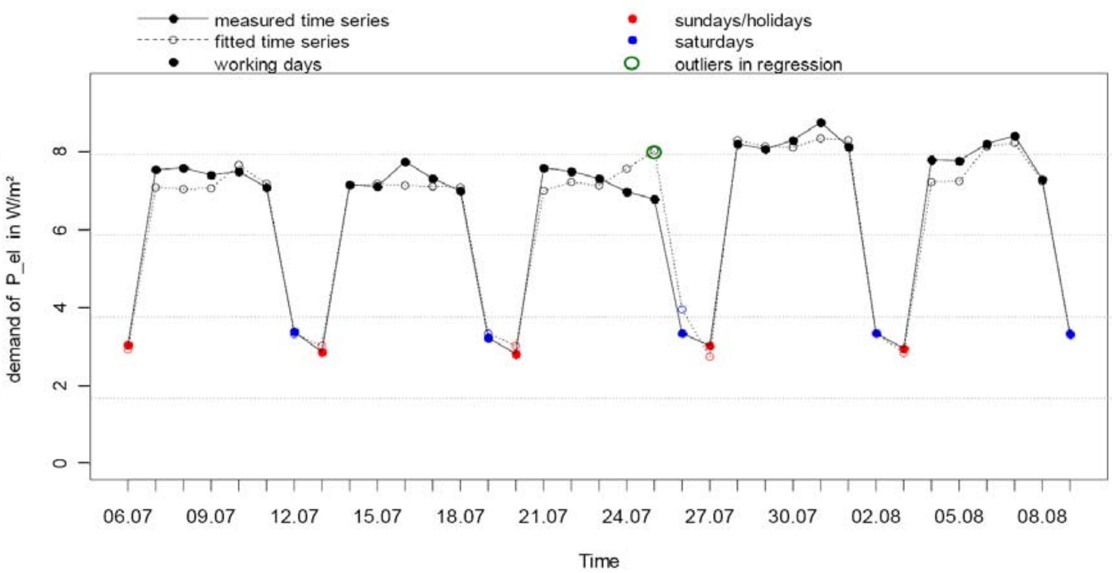

Figure 9: Result of Validation for electricity demand showing predicted values and measured values.

difference between to days ΔT_i. They were removed because they did not contain any statistically relevant information to explain the electricity demand. The office spaces in which the measurements where acquired do not represent a decent cooling profile. This illustrates that a good model can also be developed when explanatory variables are left out that normally contain information regarding the physical relationship between energy demand and temperature fluctuations. This fact is relevant for models which monitor base loads because they do not use temperature variables. They are purely statistical and mainly depend on the information of a building's occupancy status which is provided by the day-types. Even though these models are different from heating or cooling demand models their predictions are very close to the measured values, too.

CONCLUSION

The paper presented a model-based approach to monitor the energy demand in commercial buildings. The method does not rely on having extended knowledge of the HVAC equipment and the physical properties of the building. Collecting the data is the only time-consuming process.

The findings so far are very promising. The day-typing process is very reliable in determining the correct day-types during the Validation stage. The standardization of the features needs further attention and it is likely that different dissimilarity or similarity measured will be tested.

The regression models all have high adjusted R^2 values. The fitted values are very close to the measured values and relevant unusual days have been detected.

Altogether it can be said that the above process shows a lot of potential for the development of an automated, reliable and efficient tool for real-time monitoring and fault detection in buildings.

REFERENCES

[1] Seem, J.E. 2005. Pattern recognition algorithm for determining days of the week with similar energy consumption profiles. *Energy and Buildings 37* 127-139.

[2] Duda, R. O., Hart, P. E.; *Pattern classification and scene analysis*; John Wiley & Sons; 1973