CLUSTERING-BASED PREDICTION OF RESIDENTIAL ELECTRICITY CONSUMPTION

Adam Buruzs[1]*, Jan Kurzidim [1], Miloš Šipetić [2], Moritz Wagner [1]

[1]Austrian Institute of Technology, Sustainable Thermal Energy Systems, Vienna, Austria
[2]Austrian Institute of Technology, Digital resilient cities and regions, Vienna, Austria

*Corresponding Author: adam.buruzs@ait.ac.at

ABSTRACT

The EU-funded project HESTIA (Holistic demand response services for European residential communities) is striving to develop modern ICT tools for the next generation of Demand Response (DR) services for residential consumers and prosumers. For the electric load monitoring, two demo sites were set up in the Netherlands and Italy. For the energy providers, grouping/clustering of households with respect to their electric consumption profile is of great importance. It can help them to predict electricity usage, find faulty devices, and help to adjust Time of Use tariffs. Electric load forecasting has financial benefits when trading on electricity markets and offers advantages for reliable operation of electricity networks. Our present work focuses on the clustering, prediction and subsequent analysis of household electricity load profiles collected within the HESTIA project.

In 30 households both on the Italian and Dutch demo sites, the electric grid import/export data from smart meters were recorded over a one-year period with a resolution of 15 and 5 minutes respectively. The households were equipped with PV panels and home batteries, and their electricity production were monitored. The total household consumption was calculated from the grid, PV and battery values. The data quality issues required considerable efforts and preliminary analysis.

Clustering the timeseries of single household daily consumption was performed by standard clustering methods such as K-Means. This was deployed as a software service running once a day on an AIT server on real-time updated monitoring data. The clustering training is performed once a month. Many clustering methods such as K-Means involve non-deterministic steps, which makes the matching of the clusters between subsequent runs an important task. We have solved the cluster matching with the Hungarian method, where we converted the task to a linear sum assignment problem. This enables to map cluster labels, and track cluster assignments across subsequent training runs.

For predicting the cluster labels, we engineered a set of features from aggregated daily time series data, intraday data, measured and forecasted weather data, as well as calendar indicators. The cluster prediction task is interpreted as a classification problem, where the classifier is trained on historical data. Various machine learning algorithms (decision tree-based classifiers, support vector machines, neural networks) were tested for classification performance. As the input data before the clustering was normalized, the centroid of the predicted cluster represents a forecast for the shape of the next day consumption. The accuracy of this classification-based 24 h ahead electrical load curve prediction is evaluated in the presented study.

1 INTRODUCTION

Increasing the share of volatile renewables in national electricity mix includes greater flexibility challenges for electricity service operators. The ongoing electrification of space heating in buildings, as well as the increasing usage of electric vehicles increases the electrical power demand. Supply-side flexibility solutions require significant investment, but demand-side management programs have potential to deliver network flexibility at a lower cost (Miri and McPherson 2024). Demand response programs provide incentives for customers to reduce electricity consumption in specific hours of peak

https://doi.org/10.52202/077185-0062

demand and low renewable power production. Time-of-use Pricing, critical peak pricing, and real time prices are the typical financial incentives (Srivastava, Van Passel, and Laes 2018). Grouping the customers in clusters of similar usage patterns can help the utilities to design those tariff incentives. The residential consumer sector has great number of individuals with relatively small load per individual compared to industrial consumers, making aggregation of individual consumers an important analytical task. Clustering can help to better understand consumer behavior and estimate effects and applicability of time of use demand response schemes. Variability in consumption profile is also an important factor, as noted by (Wang et al. 2016), customers with lower variability are more susceptible for direct load control (Miri and McPherson 2024), whereas customers with higher variability are more suitable for price signals based demand response programs. Clustering of electrical power data is also used to detect anomalies and outliers (Capozzoli, Lauro, and Khan 2015) in building energy consumption, that can be an indication of faults.

In the EU project HESTIA residential Demand Response services are implemented and evaluated from technical and social perspectives. Demonstration households in Netherlands and Italy are equipped with smart meters and sensors for monitoring electricity consumption, PV production, home batteries charging, and smart plugs measuring device level consumptions. Based on the measurement data predictive analytics will provide forecasts of the electricity demand, and the HESTIA user profiling service enables classifying households into groups based on the similarity of their consumption patterns. This user profiling service is the subject of the present paper.

Various clustering methods were applied for load profiling i.e., grouping of daily electricity load timeseries into clusters. The applied methods include k-means, hierarchical and fuzzy clustering, and self organizing maps (Chicco 2012). The most widely used k-means clustering for electric load clustering and the effect of the number of clusters on the performance indicators (
Cluster Dispersion, Scatter Index, Davies-Bouldin index etc.) are well described in the study of (G. Chicco, Napoli, and Piglione 2006).

Many studies have been published about electrical load prediction (Runge and Zmeureanu 2019). The forecasting methods can be categorized with respect to the number of time steps predicted by the predictor. Single step ahead forecasts predict only one time step in advance, while multistep methods predict multiple time steps in advance. As is clearly visible from Table 2. and 3. of the Review paper of (Runge and Zmeureanu 2019), multistep prediction is more difficult, and usually results in higher error rates. The longer the prediction horizon, the lower is the prediction accuracy. In the present paper we target the multistep method with 24 h horizon, which is the most difficult task in the literature. Among the studies dealing with day ahead predictions the following ones deserve attention: (Chou and Tran 2018) used an ensemble of classical Machine Learning methods (SVMs, Regression Trees, and Artificial Neural Networks) to forecast on a 24 h time horizon the 15 minutes electrical loads, and obtained that a proposed hybrid method increases the prediction accuracy (The prediction-actual value correlation increased by 0.25). (La Tona, Luna, and Di Piazza 2023) used LSTM (Long Short-Term Memory) based Encoder-decoder deep Learning model for 24h ahead predictions. They have used the IHEPC dataset which consist of the electrical meter data of one individual household. (Chaouch 2014) used wavelet transformation based prediction of single household electrical load profiles using a clustering into three clusters as a preprocessing step to increase the prediction accuracy. He has performed day ahead prediction with half an hour resolution. (Teeraratkul, O'Neill, and Lall 2018) used Dynamic Time Warping to cluster daily load curves, and fitted Markov chain transition probabilities to predict the next day load curve based on the predicted representative load curve. In a recent study Gasparin et al (Gasparin, Lukovic, and Alippi 2022) systematically tested and compared multiple deep learning based models (such as Elmann Recurrent Neural Networks, Gated Recurrent Units, LSTM, Deep Convolutional Networks, seq2seq model, etc.) on the electrical load prediction task. The obtained results show relatively similar accuracy for the various models, but the variation of the underlying dataset has much higher effect on the normalized accuracy measure calculated by the regression R^2 coefficient. The aggregated dataset (GEFCom2014), which was collected as a sum of several different smart meter readings was much more regular than the individual household data (IHEPC), which resulted in a much higher prediction accuracy (R^2 of 0.8-0.9 for aggregated data vs. 0.1-0.28 for individual household data).

https://doi.org/10.52202/077185-0062

As was mentioned in the review paper of Kuster (Kuster, Rezgui, and Mourshed 2017), it is difficult to compare prediction accuracy between different research papers, as the underlying dataset and even the performance indicators used to measure the prediction accuracy (Mean Absolute Error, Mean Squared Error, R-squared, Mean Absolute Percentage Error (MAPE), etc.) show great variability among the studies.
The novelty of our contribution is the test of a simple clustering based approach, which starts with a k-means clustering of daily timeseries. Subsequently a prediction model is only trained for cluster label prediction, which is just a classification task (in contrary to the regression task what is usually targeted in the literature). The load curve prediction in our case is only a byproduct of the label prediction. Additionally, we also present a solution for cluster label matching, which is an important task in real applications to match cluster groups after retraining of the clustering. The current study also discusses the specific data preparation process performed during the project.

2 HESTIA DATASET

Within Hestia project 2 demo sites were set up to collect realtime data from smart sensors, one in Italy and the other in Netherlands. The current study used the dataset from the demonstration site in Italy. 30 Households located in Berchidda were equipped with IoT sensors and hubs of Nesosnet company. The smart meters are communicating with the MQTT broker via the MQTT protocol, and this MQTT server is forwarding the collected data to a central Influxdb database, where all HESTIA data is stored. Detailed information can be found in deliverables 5.1 and 5.2 on the HESTIA website (AdminHestia 2024). Most households have been also equipped with PV panels (with capacity of 0-7 kWp) and LiFePO home batteries, where the inverters are sending produced/consumed power data to the database via the MQTT broker. Furthermore, weather data (temperature, humidity, solar radiation, precipitation etc.), as well as the next day prediction of these values were collected from the weatherbit services (weatherbit.io).
The smart meters were measuring the total electricity consumption of the households inclusive the battery (dis)charging and PV production. As in this study we are only focusing of the electrical consumption of the households, we have to factor out this data, so for the household electric consumption the following calculation formula was used:

$$\text{household\_load} = \text{grid\_import} - \text{grid\_export} + \text{PV\_production} - \text{battery\_power} \quad (1)$$

This calculation formula assumes that all the 4 terms on the right side are correct (no data errors). As the PV and battery inverters were not synchronized with the smart meter, the collected data has slightly different timestamps. To match those database tables a 15-minute resampling of the data was applied. Missing values of the PV and battery power data were filled with zeros. Resulting from various synchronization issues some load values (less than 1 %) calculated by Eq.1 became negative. These values are set to zero.

3 CLUSTERING AND CLUSTER MATCHING

3.1 Clustering

Clustering of electricity load profiles is extensively covered in the state of the art literature, for example the study of Chicco et al from 2006 (G. Chicco, Napoli, and Piglione 2006) compares the performance of k-means, fuzzy k-means, self-organizing maps (SOM) and hierarchical clustering. The number of clusters parameter must be entered as an initial parameter to the k-means clustering algorithm. As it is described by Chicco et.al there are some indicators - like the knee of the cluster dispersion, or the scatter index - used in the literature to find the “best” number of clusters, but they often result in relatively high number of clusters (e.g., about 40), which is not feasible for the service operator to keep track of. Therefore, it’s often more practical to let the service operator (utilities) determine the number of customer groups to be distinguished.

https://doi.org/10.52202/077185-0062

The electrical consumption load data was treated as a timeseries with 15 min resolution, and for each household every day (0-24 h) of power consumption values was treated as one datapoint i.e., a vector of 4*24 = 96 values. This data was normalized so that each such vector has unique mean value (each household day vector has the same sum). With this normalization, for example, if one household has exactly double the consumption compared to another household with the same shape, then they correspond to the same point in the feature space, this implies this clustering is invariant when scaling the consumption by a factor. These 96 element long arrays of household-day consumptions were fed into the k-means clustering algorithm, without dimension reduction step. Negative values and outliers were removed.

In Figure 1 the resulting cluster centers of the clustering with 8 clusters is visible. The difference between various consumption patterns is also easy to interpret, cluster 2 and 5 having high consumption at evening, and cluster 3 and 1 rather high consumption in the morning hours. The most frequent cluster 4 has no well-defined consumption peak.

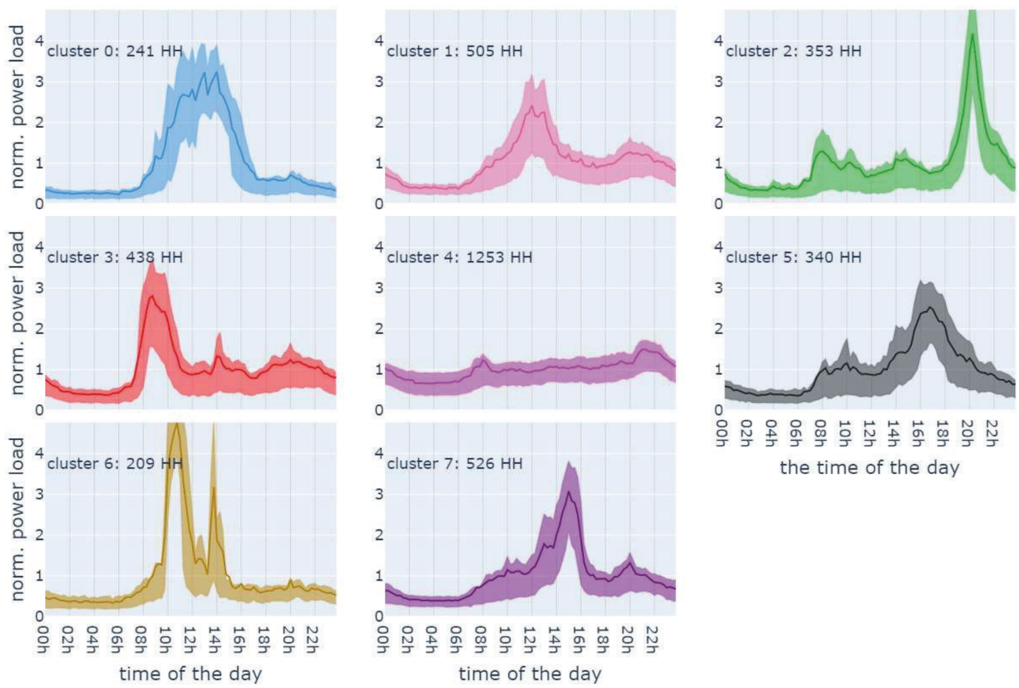

Figure 1 The normalized daily electric load profiles of the resulting clusters.

3.2 Cluster matching

The k-means clustering algorithm starts with a randomly assigned cluster centers. These cluster centers are iteratively updated, which results in a cluster label assignment, that is dependent on the initial random center selection. Using different random seeds for initialization can lead to different cluster assignments even on the same dataset by running the k-means clustering multiple times. If the dataset is growing/changing and the cluster assignments need to be updated (as was the case by the HESTIA user-profiling services) the clustering algorithm must be recalculated, which can lead to inconsistent cluster assignments. The presented cluster matching method is solving this problem: Having two set of cluster center assignments (results of two independent k-means clustering) with same number of clusters, calculate the matching between the two sets, i.e., which clusters from clustering B corresponds to clusters of A? We solve this problem with the help of the Hungarian Method (Kuhn 1955), which was originally invented for the assignment problem, but also successfully used to cluster matching (Kume and Walker 2021). The method finds the permutation of cluster centers of clustering run B, so that the sum of the distances between centers of A, and the corresponding clusters of B is minimal. For the matching we use the linear sum assignment implementation of the scipy Python package.

https://doi.org/10.52202/077185-0062

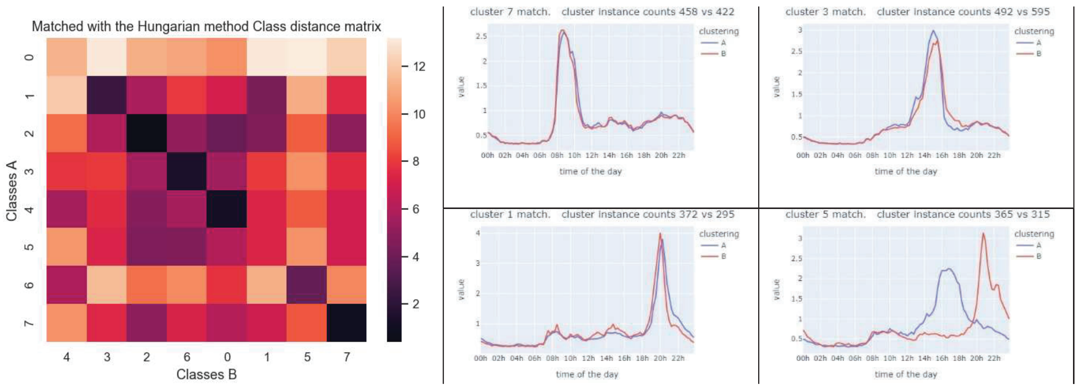

Figure 2 Cluster matching example. Left: the distance matrix between the matched clusters. Right: examples of matched cluster centroid load profiles.

As an example, we generated two k-means clustering runs with 8 clusters with two different random seeds. In Figure 2 we show the distance matrix between clustering A, and permutation of clusters B. The permutation is the result of the Hungarian algorithm applied to the distance matrix. Ideally, we have low numbers in the diagonal, that means the cluster centers are close to each other in the feature space. On the right we showed the matched centroids. As we see from the images some cluster centers match very well, but some are different (right bottom profiles). In the heatmap of the left is already visible that cluster 5 of clustering A is not matching well to cluster 1 of clustering B, but still that match leads to the best overall matching.

4 CLUSTER LABEL AND LOAD PREDICTION

Electric load forecasting offers the advantage of better generation capacity management and scheduling for the electric utilities companies, forecasting also contributes to maintaining the grid stability, especially if volatile renewables have big share in the total power production. Forecasts should also play significant role in demand response programs, where electricity consumers are incentivized to modify and or shift their peak electricity demand during specific periods (Srivastava, Van Passel, and Laes 2018).

As a first step we solved the task of cluster label prediction: using data available today, and given our trained clustering with N clusters, what will be the cluster label of a given household on the next day? We have treated this problem as a classification problem of machine learning. We can generate a set of features from the data of the current and previous days and train a classifier to predict the class (cluster label) of the next day. The set of features engineered manually and consist of five different types of features:

- Features based on daily aggregated consumption data (daily one value per timeseries) (example: total daily consumption/ average consumption over last 7 days)
- Features based on intraday data (15 minutes resolution) (example: the hour of maximal consumption)
- Measured weather data (example: daily maximum temperature)
- Weather forecast data (example: UV radiation prediction for the next day)
- Calendar indicator (example: day of the week)

For every household around 40 numerical features are generated for each day, each belonging to one of the above categories. Additionally, the cluster label of the actual day was also used as input. Based on the calculated feature table we predict the cluster label of the next day (class). The input features were scaled by shifting the median to zero and scaled by the 10-90 % quantile range. This scaling is robust to outliers and anomalies that has less than 10 % frequency. We have trained and compared multiple basic machine learning models for classification, including Support Vector Machines, Decision Tree, Random Forest, Gradient Boosted Tree, Nearest Neighbors classifier, and Multi Layer Perceptor Neural Network. Sklearn (Pedregosa et al. 2011) Python package was used for the implementation.

https://doi.org/10.52202/077185-0062

The basic algorithms shown in Table 1 were tested on the data from jan 2023 – dec 2023, with a target k-means clustering labels of 8 different clusters. Hyperparameter combinations shown in the second column were tested with a 5-fold cross validation. The mean cross validation score of the best instance is shown in the last column (F1 score). The parameters of the best instance are marked bold in the second column. The gradient boosting tree ensemble classifier (Friedman 2001) results in the highest F1 score. This model gives an accuracy (rate of correct class label prediction) of 49 %. The Multilayer Perceptor Neural Network and the Random Forest predictor are also resulting higher scores, as these models fit many trainable parameters, and can discover complex dependency patterns. The k-nearest neighbors model is the simplest model with only the number of neighbors as parameter, that explains the low F1 score. Random forest is based on single decision trees but using bagging (bootstrap aggregation) to train multiple (in our case 100) decision trees on random subset of the data, and uses these ensemble of decision trees for an aggregated decision (Jiawei 2012). This method is developed to improve classification accuracy of decision trees, which we can also confirm with our result. Gradient Boosting also belongs to the class of ensemble methods, where in each step an additional decision tree (a weak learner) is fitted to improve the log-loss function of regression/classification using gradient descent optimization.

Table 1 Results of hyperparameter tuning of the basic Machine Learning algorithms, for the next day cluster label prediction task.

Classifier Algorithm	**Tested Hyperparameters**	**F1 score**
K nearest neighbors	*n_neighbors*: [5, **20**, 100]	0.283
Decision Tree	min_samples_split: [10, 20, **50**], max_depth: [3, 5, 8, **12**], criterion: [**'gini'**, 'entropy', 'log_loss']	0.330
Random Forest	n_estimators: [5, 20, **100**] max_depth: [2, 4, **6**] min_samples_split: [**10**, 20, 50] criterion: [**'gini'**, 'entropy', 'log_loss']	0.346
Gradient Boosting	min_samples_leaf: [10,50, **100**], max_depth:[**2**,3,4, 5], learning_rate: [**0.1**,0.3], n_estimators: [20, **50**, 100]	**0.376**
Support Vector Machine	kernel: ['**rbf**', 'poly'], C: [**1.0**, 0.2]	0.337
Multi Layer Perceptor	solver: ["lbfgs", "**sgd**", "adam"], activation: [**'tanh'**, 'relu', 'logistic'], hidden_layer_sizes: [(20,), (100,), (30, 20,), (20,30,20), **(50,50)**], n_iter_no_change: [10, **40**]	0.349

By the gradient boosting classifier, the feature importance scores can be queried for each feature. This is the rate by which a given feature reduces the Gini impurity in total, normalized so that the feature importance scores sum to 1 for all features. The top 3 features were:

- *minRate* = (the lowest 3 hour consumption during the day) / (average daily consumption) (score = 0.31)
- Last Cluster label (score = 0.16)
- Second moment (*mom2*) of the hourly consumptions (score = 0.15)

Where the second moment is defined as

$$mom2 = \frac{1}{x_{av}} \sum_{h=1}^{24} (x(h) - x_{av}) h^2 \qquad (2)$$

With the sum is running over the hours *(h)*. *x(h)* is the actual electrical power consumption over the given hour, and x_{av} is the average consumption over the whole day. The confusion matrix of the class label prediction is shown on Figure 3.

https://doi.org/10.52202/077185-0062

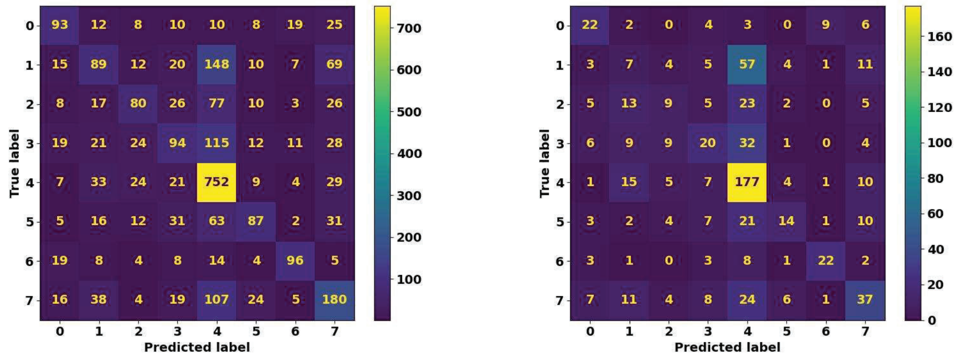

Figure 3 Confusion matrix of the cluster label prediction with GBC classifier. Left on the training set, right on the test set.

With the predicted clusters we can construct a prediction of the power consumption profile with 15 minutes resolution and 24h forecast horizon. As a simplest approximation our prediction of the next day profile can be the cluster mean profile of the predicted cluster. One example of such prediction is shown on Figure 4. As the clustering is applied to the normalized daily consumption data, the predicted profile is also normalized (the average normalized power is one.) The quality of the prediction can be measured by the mean squared error (MSE), which is calculated from the difference between the Predicted ($P^{i}{}_{pred}$) and (ground) true $P^{i}{}_{tru}$ power load over all N datapoints.

$$MSE = \frac{1}{N}\sum_{i=1}^{N}\left(P^{i}{}_{pred} - P^{i}{}_{true}\right)^2 \tag{3}$$

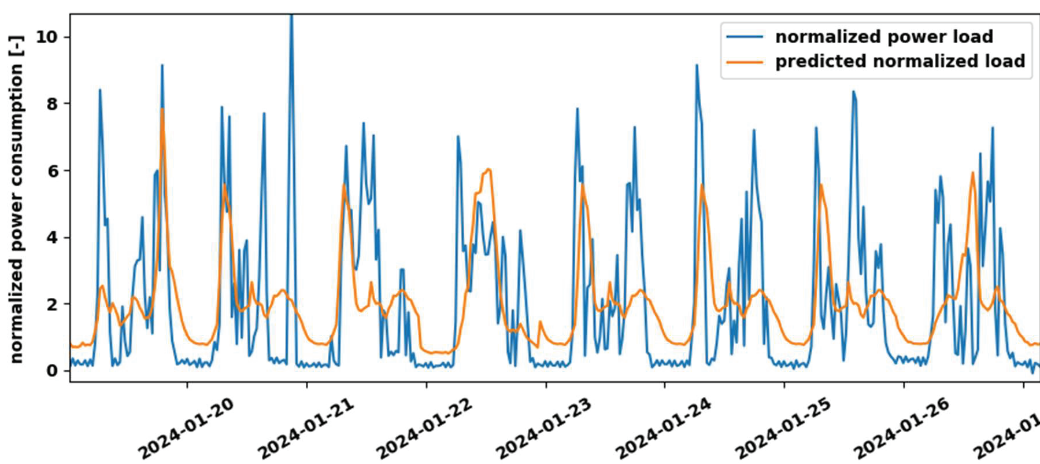

Figure 4 Predicted power profile vs the actual profile

To better estimate the predicting power of our forecaster we can calculate a reference baseline: As a baseline we can consider the mean-predictor which predicts the mean (1) normalized power for each 15-minute interval point. The baseline predictor results in MSE of 1.14 (see column MSE_baseline_C1 on Table 2). Our clustering-based predictor for 8 clusters results in 0.95 (80 % training data) and 0.98 (test data), so it shows 14 % better prediction error than the baseline. We have also calculated what error would we get if we could perfectly predict the next day cluster (looking in the future). For 8 clusters prediction we would have a MSE of 0.74 with this hypothetical ideal cluster-based predictor (see MSE_ideal column in Table 2). This is the theoretical upper limit of accuracy that can be reached with cluster label-based profile predictor. The increased error compared to this ideal predictor is the effect of inaccuracy of cluster label predictions. If we increase the number of clusters, the ideal predictor gets better, as the cluster centers correspond to smaller number of cluster instances, and the cluster center has smaller distance to the instances. However, the cluster label prediction is more difficult with

https://doi.org/10.52202/077185-0062

more clusters (bigger confusion matrix), that's why the real profile predictor is not improving with varying number of clusters (see Table 2: MSE_GBS_predictor is the total MSE of the GBC predictor for all data, and in column "MSE Train/Test" is the MSE separately for the train and test set). The mean absolute error (MAE) of the normalized power prediction is 0.60. This is only slightly weaker as the results of (La Tona, Luna, and Di Piazza 2023), who trained an LSTM neural network to forecast electrical load profiles, and obtained a MAE of 0.5-0.55. However, we have to stress, that the two results are not directly comparable, as they were tested on different datasets. As our clustering-based method optimized the machine learning algorithm for the best class label prediction, it is plausible, that this does not result in the optimal load profile prediction. As discussed in the review of Runge (Runge and Zmeureanu 2019), the 24h multistep ahead forecast for 15 minutes energy consumption timeseries is a relatively difficult challenge compared to single step forecasting, and shorter horizon multistep forecasting tasks, leading to higher mean absolute percentage error (MAPE) of around 36-42 %.

Table 2 Mean Squared Error (MSE) results with the clustering based normalized load profile predictor with various cluster counts.

N cluster	MSE_baseline_c1	MSE_ideal	MSE_GBC predictor	MSE Train/Test
5	1.14	0.78	0.97	0.98/0.96
8	1.14	0.74	0.96	0.95/0.98
12	1.14	0.70	0.97	0.97/0.97

5 CONCLUSION AND OUTLOOK

In the present work we describe a k-means based clustering of residential electrical load profiles, based on the dataset which was generated during the Hestia EU project. The results show clear identification of electrical usage patterns like morning and evening consumers. We also describe how the identified clusters can be matched over subsequent clustering runs. We have also trained machine learning models, for predicting the next day cluster labels based on the historical load, weather observations and weather forecasts. We show how these cluster label predictions can be used for timeseries profile shape predictions.

We recognize the potential to enhance prediction accuracy by augmenting the applied features with additional information from the electrical load timeseries. Specifically, application of deep learning models has the potential to extract meaningful representations as features discovered by neural networks. Still, one has to be aware the limitations of prediction accuracy as the schedule and calendar of the residents are not available as input data due to privacy concerns, however many people tend to have individual daily agenda, which has great influence on their household electrical consumption.

Even if the prediction of load curves of single households shows high error rates, the aggregated load curves of multiple households expected to have lower variability (by smoothing out individual behavior) and also reported in the literature. Therefore, the application of the presented single household predictions on aggregated load curves is a promising task for future projects.

6 REFERENCES

AdminHestia. 2024. "Project Deliverables." *HESTIA* (blog). 2024. https://hestia-eu.com/resources/project-deliverables/.

Capozzoli, Alfonso, Fiorella Lauro, and Imran Khan. 2015. "Fault Detection Analysis Using Data Mining Techniques for a Cluster of Smart Office Buildings." *Expert Systems with Applications* 42 (9): 4324–38. https://doi.org/10.1016/j.eswa.2015.01.010.

Chaouch, Mohamed. 2014. "Clustering-Based Improvement of Nonparametric Functional Time Series Forecasting: Application to Intra-Day Household-Level Load Curves." *IEEE Transactions on Smart Grid* 5 (1): 411–19. https://doi.org/10.1109/TSG.2013.2277171.

Chicco, G., R. Napoli, and F. Piglione. 2006. "Comparisons among Clustering Techniques for Electricity Customer Classification." *IEEE Transactions on Power Systems* 21 (2): 933–40. https://doi.org/10.1109/TPWRS.2006.873122.

https://doi.org/10.52202/077185-0062

Chicco, G. 2012. "Overview and Performance Assessment of the Clustering Methods for Electrical Load Pattern Grouping." *Energy* 42 (1): 68–80. https://doi.org/10.1016/j.energy.2011.12.031.
Chou, Jui-Sheng, and Duc-Son Tran. 2018. "Forecasting Energy Consumption Time Series Using Machine Learning Techniques Based on Usage Patterns of Residential Householders." *Energy* 165 (December):709–26. https://doi.org/10.1016/j.energy.2018.09.144.
Friedman, Jerome H. 2001. "Greedy Function Approximation: A Gradient Boosting Machine." *The Annals of Statistics* 29 (5): 1189–1232.
Gasparin, Alberto, Slobodan Lukovic, and Cesare Alippi. 2022. "Deep Learning for Time Series Forecasting: The Electric Load Case." *CAAI TRANSACTIONS ON INTELLIGENCE TECHNOLOGY* 7 (1): 1–25. https://doi.org/10.1049/cit2.12060.
Jiawei, Han. 2012. *Data Mining*. Elsevier. https://doi.org/10.1016/C2009-0-61819-5.
Kuhn, H. W. 1955. "The Hungarian Method for the Assignment Problem." *Naval Research Logistics Quarterly* 2 (1–2): 83–97. https://doi.org/10.1002/nav.3800020109.
Kume, Alfred, and Stephen G. Walker. 2021. "The Utility of Clusters and a Hungarian Clustering Algorithm." *PLoS ONE* 16 (8): e0255174. https://doi.org/10.1371/journal.pone.0255174.
Kuster, Corentin, Yacine Rezgui, and Monjur Mourshed. 2017. "Electrical Load Forecasting Models: A Critical Systematic Review." *Sustainable Cities and Society* 35:257–70. https://doi.org/10.1016/j.scs.2017.08.009.
La Tona, G., M. Luna, and M. C. Di Piazza. 2023. "Day-Ahead Forecasting of Residential Electric Power Consumption for Energy Management Using Long Short-Term Memory Encoder–Decoder Model." *Mathematics and Computers in Simulation*, July. https://doi.org/10.1016/j.matcom.2023.06.017.
Miri, Mohammad, and Madeleine McPherson. 2024. "Demand Response Programs: Comparing Price Signals and Direct Load Control." *Energy* 288 (February):129673. https://doi.org/10.1016/j.energy.2023.129673.
Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, et al. 2011. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research* 12 (85): 2825–30.
Runge, Jason, and Radu Zmeureanu. 2019. "Forecasting Energy Use in Buildings Using Artificial Neural Networks: A Review." *Energies* 12 (17): 3254. https://doi.org/10.3390/en12173254.
Srivastava, Aman, Steven Van Passel, and Erik Laes. 2018. "Assessing the Success of Electricity Demand Response Programs: A Meta-Analysis." *Energy Research & Social Science* 40 (June):110–17. https://doi.org/10.1016/j.erss.2017.12.005.
Teeraratkul, Thanchanok, Daniel O'Neill, and Sanjay Lall. 2018. "Shape-Based Approach to Household Electric Load Curve Clustering and Prediction." *IEEE Transactions on Smart Grid* 9 (5): 5196–5206. https://doi.org/10.1109/TSG.2017.2683461.
Wang, Yi, Qixin Chen, Chongqing Kang, and Qing Xia. 2016. "Clustering of Electricity Consumption Behavior Dynamics Toward Big Data Applications." *IEEE Transactions on Smart Grid* 7 (September):2437–47. https://doi.org/10.1109/TSG.2016.2548565.

ACKNOWLEDGEMENT

The presented research is financed by the European Union's H2020 programme under Grant Agreement No. 957823.

https://doi.org/10.52202/077185-0062